VISUAL COGNITION, 2000, *7* (1/2/3), 175–190

What's Scene and Not Seen: Influences of Movement and Task Upon What We See

Guy Wallis

Department of Human Movement Studies, Queensland University, St Lucia, Australia

Heinrich Bülthoff

Max-Planck Institute for Biological Cybernetics, Tübingen, Germany

Studies concerning the processing of natural scenes using eye movement equipment have revealed that observers retain surprisingly little information from one fixation to the next. Other studies, in which fixation remained constant while elements within the scene were changed, have shown that, even without refixation, objects within a scene are surprisingly poorly represented. Although this effect has been studied in some detail in static scenes, there has been relatively little work on scenes as we would normally experience them, namely dynamic and ever changing. This paper describes a comparable form of change blindness in dynamic scenes, in which detection is performed in the presence of simulated observer motion. The study also describes how change blindness is affected by the manner in which the observer interacts with the environment, by comparing detection performance of an observer as the passenger or driver of a car. The experiments show that observer motion reduces the detection of orientation and location changes, and that the task of driving causes a concentration of object analysis on or near the line of motion, relative to passive viewing of the same scene.

INTRODUCTION

One of the most remarkable abilities of the human visual system is its capacity to reliably interpret complex natural scenes, not least when faced with surroundings in a state of flux. The details of how we achieve this apparently

---

Please address all correspondence to G. Wallis, Department of Human Movement Studies, Queensland University, St. Lucia 4072, Queensland, Australia. Email: gwallis@hms.uq.edu.au

We are grateful to William Hayward, to the editor, and the two referees for helpful comments and suggestions. Guy Wallis was supported by a grant from the Motor Accident Insurance Commission and a Max-Planck Society Research Fellowship.

accurate interpretation remain under debate. However, recent research has succeeded in demonstrating that, contrary to our subjective experience, scenes are not represented in full and complete detail, but rather in a much more sparse and approximate manner. Indeed it seems that detailed analysis of the constituent parts of a scene is not instantaneous, and can only be achieved in a more piece-wise, serial manner (see Simons & Levin, 1997, for review). This raises interesting questions about what elements of a scene are encoded in detail, and under what circumstances. Ultimately, one would like to understand how we can successfully interact with a scene despite the actual vagueness with which many of the objects, or their attributes (colour, location, form, etc.) are encoded.

One of the most direct means of testing whether a subject has processed information about a particular part of a scene is to change the associated information, and then record whether the subject is sensitive to this change or not. The problem with this approach is that any local change within the scene will produce a local, apparent motion signal. Humans are extremely sensitive to motion in all parts of their visual field, and a motion signal of this sort will draw attention to the region of change, obscuring the question of whether or not the object was the focus of attention before the change took place. At least one ingenious solution to this problem is to make changes contingent upon saccadic eye-movements. During a saccadic eye-movement our sensitivity to motion (among other attributes) is greatly reduced, and thus the apparent motion signal associated with making a change is suppressed. Studies using this technique have shown that astonishingly large changes can be made to the composition of individual static scenes without them being immediately obvious to passive observers (Grimes, 1996; Irwin, 1991; McConkie & Currie, 1996; McConkie & Zola, 1979). Unfortunately, saccade contingent changes are technically very difficult to produce, requiring specialised equipment and, until recently, some degree of head restraint. Fortunately, other researchers have shown that alternative techniques can be used to mask apparent motion signals. These techniques include introducing a grey, blanking image between changes (Rensink, O'Regan, & Clark, 1997, Rensink, O'Regan, & Clark, this issue), making changes contingent upon eye-blinks (O'Regan, Deubel, Clark, & Rensink, this issue), and making lateral shifts of the entire image between changes (Blackmore, Brelstaff, Nelson, & Troscianko, 1995).

A common element of all of the work just described is that it focuses on the analysis of static scenes. In everyday life of course, we are actually exposed to a continuously changing retinal image. This might be due to head, eye, or body movements, or indeed the motion of objects within the scene. It seems plausible that the associated retinal flux will affect the characteristics of our insensitivity to changes. One situation in which changes to dynamic scenes have been made (albeit inadvertently) is in the film industry. For film-makers it is important that the continuity editor does his or her job properly, to prevent changes from occurring when cutting between shots of a scene. Although their work is important, Hochberg (1986) noted that in practice many changes in films go unnoticed, from which he argued that the representation of dynamic scenes is incomplete, much as it is for static scenes. Levin and Simons (1997) have looked into this more systematically, by making short films in which changes are deliberately made between cuts. Their work revealed a remarkable insensitivity to changes such as to the attire or even identity of persons appearing in the film.

Although film editing brings us somewhat closer to the appearance of real scenes, the abrupt changes in viewpoint typical of film editing make it difficult to assess our sensitivity to changes in the more natural setting of a smooth, continuous environment. Recent work has therefore sought to describe change blindness in real, dynamic scenes (Simons & Levin, 1997). In this paper we too aim to measure change detection in a dynamic scene and discuss an approach which combines several of the techniques mentioned earlier with a scene in which the subject' s viewpoint varies smoothly and continuously.

Apart from the effects of observer motion, we also aim to characterize the role of task in change detection. We were specifically interested in the question of how the level of interactivity of a subject affects the analysis of a scene. It has already been shown that active observers build a representation that is more effective for solving certain tasks, than do passive observers. For example, subjects who control a simulated flight-path are better able to extrapolate from their current position to one in the near future (Larish & Andersen, 1995), and active subjects make more accurate heading judgements (Telford, Howard, & Ohimi, 1995). There are also many studies in the human navigation literature, which have identified improved way-finding in car drivers, in comparison with car passengers (Péruch, Vercher, & Gauthier, 1995; Wilson, Foreman, Gillett, & Stanton, 1997). There are presumably several reasons for these differences but at least one is the manner in which the subjects represent the space around them as a function of their particular task.

In satisfying both aims we develop an environment in which objects change in colour, form, and so on, during simulated observer motion. The results not only reveal regions of interest, akin to Rensink et al.' s "centers of interest", but also the systematic manner in which the boundaries of these regions depend upon task. The results also reveal how the level of subject interaction affects the level of sensitivity to certain changes.

EFFECTS OF MOTION AND TASK ON CHANGE BLINDNESS

When viewing a scene, we are continually attending to different constituent parts or objects. The choice of what we attend to is a complex mixture of context and prior experience, and this choice will, in turn, directly affect the speed with which subjects detect changes to objects within that scene (Rensink et al., 1997). Likewise, the processing time required to recognize objects is affected by an experience driven concept of the likelihood of the object appearing in a particular scene, and of its appearing in a particular position in the scene (Biederman, Mezzanotte, & Rabinowitz, 1982; Boyce & Pollatsek, 1992; Boyce, Pollatsek, & Rayner, 1989). Sofas, for example, should not appear in front of petrol stations, and should be on the ground, not in the air (Biederman et al., 1982). Such expectations also guide our eye-movements. Indeed, numerous studies on driving have reported the effect of driving experience on the frequency of eye-movements and choice of fixation targets (Mourant & Rockwell, 1972; Previc, 1990).

What is included in the current context can be more than simply the scene itself. For instance, it may also include prior information in the form of instructions to the subject (Aginsky & Tarr, this issue). In other words, the subject' s task will play an important role in deciding where to look and what to attend to. In the following experiment we seek to characterize variation in the region of attention of an active subject (driver), relative to a passive subject (passenger). In this way we seek to characterize the influence of task on where subjects choose to focus their attention. In addition, we also seek to characterize the affect of observer motion on change detection in general.

The first hurdle to answering these questions is the development of a technique capable of measuring change blindness in an interactive and controlled environment. The standard approach in scene analysis experiments has been to test detection on real scenes, in the form of videos and photographs. The main disadvantage of this approach for our study is the restricted level of subject interaction this affords. As we are particularly interested in the question of how observer interaction affects scene analysis, we would prefer to use a virtual environment instead, that is, a realistic-looking environment generated by a computer. In practice, transferral of the change blindness technique to a virtual environment carries both advantages and disadvantages. The disadvantages all centre around the question of how real the virtual environment is. Virtual environments use very simplified lighting models, the movement may be jerky at high angular velocities, and many other sensory cues, such as vestibular or proprioceptive ones, may be missing (Carr & England, 1993; Kalawsky, 1993). On the other hand, the fact that in a virtual world one can control all of the available visual cues and types of change greatly facilitates the design of complete, fully balanced psychophysical experiments.

Ultimately, if one wants to test the effect of an interactive task on change detection, then video sequences are simply not flexible enough and one has to resort to using a virtual scene. However, given the shortcomings of virtual environments listed earlier, before we can proceed to the main experiment it is important first to establish a link between detection performance in a virtual world to that in a real world. To that end, before embarking on the main experiment we first describe a pilot experiment, aimed at providing an overall comparison of performance in virtual and real environments.

EXPERIMENT I

Method

In order to present a dynamic, real world scene to subjects we modified the standard flicker technique as described by Rensink et al., (1997)—see Fig. 1(a). Our modification of the paradigm is shown in Fig. 1(b), in which each image is replaced by several frames from a video. By filming a road from a moving vehicle it was possible to give subjects the impression of smooth movement along the road. By then repeatedly filming the same stretch of road, each time under slightly differing conditions, it was possible to introduce changes to the scene.

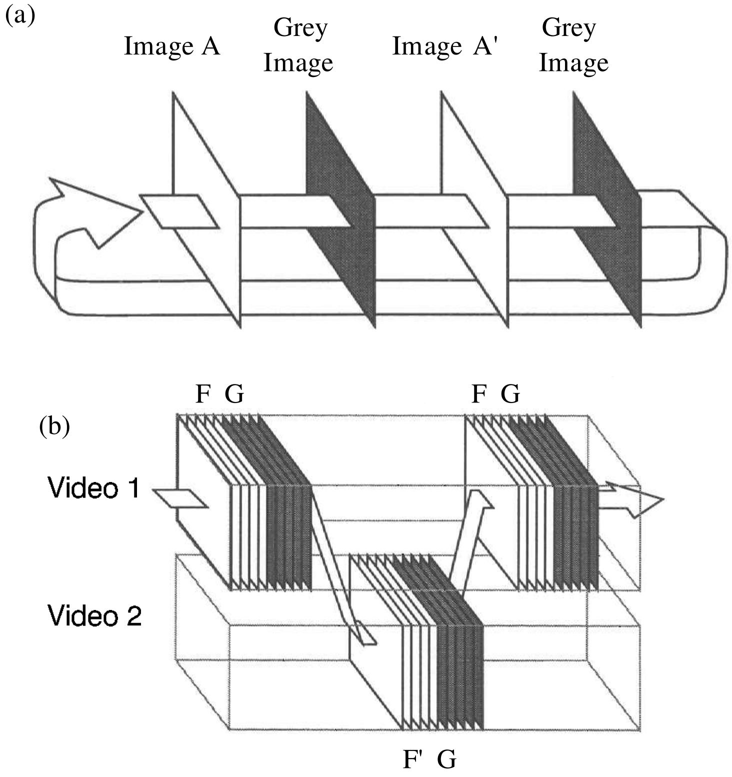

FIG. 1. (a) The standard flicker paradigm. During a trial the first image A and its altered counterpart A' are displayed repeatedly until a subject notices the difference between the two scenes. (b) Adaptation of the flicker technique to generate dynamic scenes using the mixing of frames from two video recordings of a scene.

The scene used in our investigation was a suburban street which had been closed to normal traffic, and which we filmed from a car travelling at 28 km/h—see Fig. 2. The scene contained five objects for use in the experiment (bench, small box, large box, umbrella, and ball), which were placed in a set configuration along a line perpendicular to the line of the road. The same stretch of road was filmed a total of 12 times to produce 12 videos, which differed from one another only in that the properties of two of the five objects were altered relative to the standard configuration. For example, in video 1 the small box appeared red, and in video 2 it was green, or the bench was upright in video 1 and appeared tilted backwards in video 2. Each video contained approximately 400 video frames, each equivalent to 16 sec viewing time.

In order to produce the video equivalent of the flicker images used by Rensink et al. (1997), random pairs of the 12 videos were then mixed to produce 30 new video sequences. Mixing involved recording 8 frames from video 1, 8 frames of a uniform grey image, 8 frames from video 2, a further 8 grey images, and then back to video 1, until 400 frames had been processed—see Fig. 1(b). The film was edited together in such a way that the observer was given the impression of smooth continuous motion towards and ultimately past the changing objects. In other words, the simulated motion was maintained through the blank interval too, to prevent unnatural, abrupt stopping and starting. As a result of the mixing process, from 1 to 4 changes could occur in any one of the new video sequences. Possible changes included alterations in an object' s colour, orientation, position or presence (i.e. it could be removed). In practice most videos contained 3 or 4 changes with an average of 3.6 changes.

A total of six subjects were shown each of the 30 video sequences which were displayed at standard PAL video resolution (768 × 576 pixels, subtending 30° × 22.5°) on a Silicon Graphics (Onyx Reality Engine) computer monitor. On detection of a change, subjects were initially only required to indicate this by pressing a key. At the end of each trial they were then asked to give an oral description of the changes they had seen. By saving reporting until the end of the trial, it is possible that a certain degree of reporting accuracy was lost due to

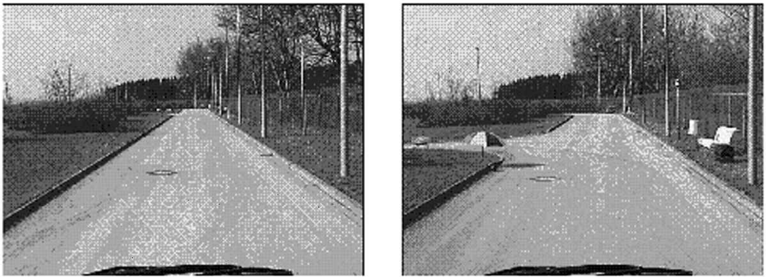

FIG. 2. Examples of frames from the video images shown to subjects.

forgetting. However, this was regarded as preferable to immediate reporting which might have interfered with the continuing search for other changes.

In a second series of tests, a further six subjects were shown a virtual model of the scene displayed in the videos. The number and types of change that were tested were exactly matched to those of the video sequences, as was the display resolution and size. Figure 3 shows the virtual environment at two instances during a typical trial. These correspond to the same points shown in the video images in Fig. 2. Note that subjects once again reported changes at the end of a trial and thus, if forgetting did occur, it is likely that it affected both the real and virtual environments equally, permitting a comparative study of performance to be made.

Results

The results displayed in Fig. 4 demonstrate that, as in previous work on static scenes, detection rates are affected by the type of change, as well as by the particular object changed. The results also reveal that performance in the real and virtual environments is very similar. That said, performance in the virtual environment is higher overall than in the real environment. The number of changes processed correctly per trial (N_C) was 2.13 for the video images, compared to 2.53 for the virtual images. Calculating N_C for all subjects this difference proved to be statistically significant, $t_{10} = 5.12$, $p < .002$. However, the fact that the improvement in detection in the virtual environment is across all objects and all change types, suggests that this improvement may well be ascribed to a general property of the images used, such as their contrast or brightness. Of more importance is the general consistency in performance across all object and change types in the two presentation paradigms.

Conclusions

In general, we can conclude that change blindness occurs in dynamic, smoothly changing environments, just as it does in static ones or between cuts in films. We can also add that the video mixing technique described represents a valid

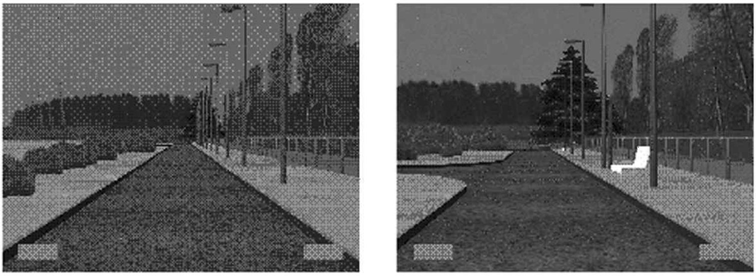

FIG. 3. Views of the 3D environment, modelled on the original videoed scene.

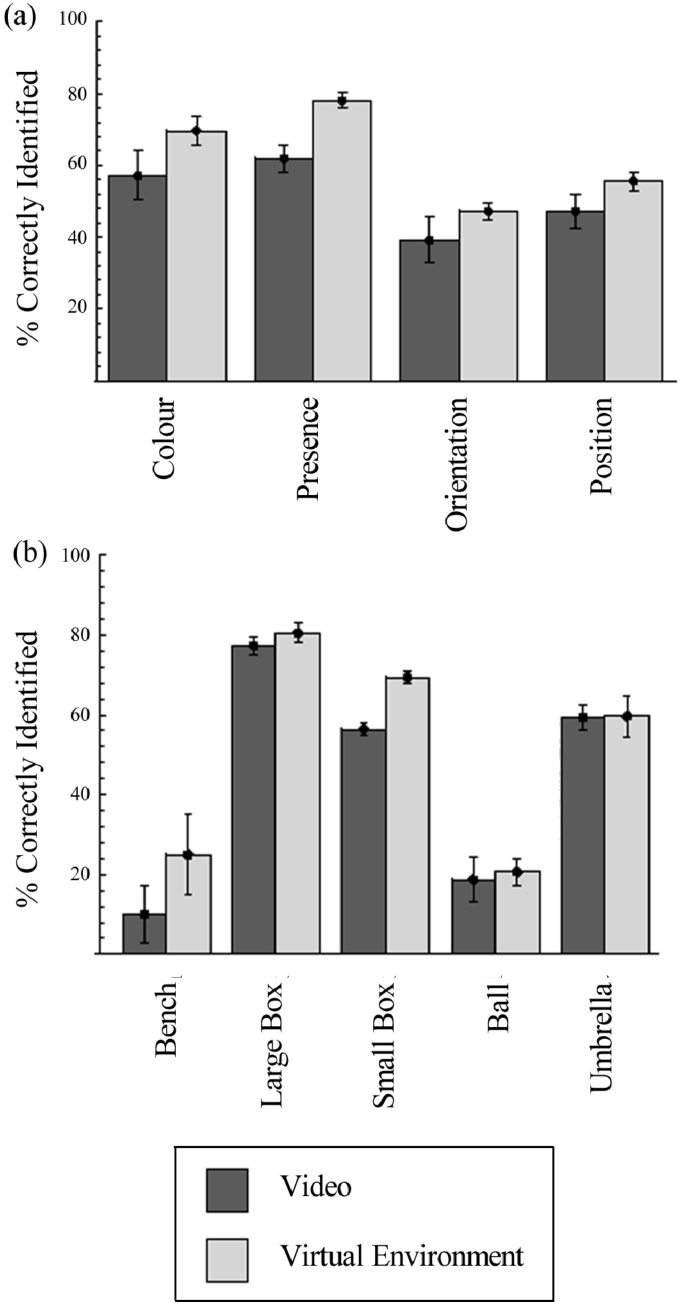

FIG. 4. (a) Percentage of changes correctly identified, arranged by type of change. (b) The same results ordered by object type. Error bars indicate *SE* of the mean in each case.

development of the standard flicker paradigm, and is suitable for testing change blindness in dynamic environments. The experiment also reveals that although possibly different in magnitude the relative detectability of changes in a virtual environment is comparable to those obtained in a real environment. This in turn serves to justify the use of a virtual environment in change blindness studies, as described in the following experiment.

EXPERIMENT II

Method

In order to investigate the main questions of this paper in more detail, a model of a 3D environment was rendered using a Silicon Graphics computer and projected onto a large, semicircular projection screen (3 m radius × 1.7 m tall, 180° × 50°). By manipulating the simulated viewpoint within the scene, subjects experienced simulated locomotion. The type of locomotion was divided into one of three categories. Subjects moved either passively through the scene (Passive), viewed the scene statically from five viewpoints (Static), or actively steered through the scene (Active)—see Fig. 5(a).

During the experiment, a total of 60 coloured cubes (1 m × 1 m × 1 m) were distributed randomly on or next to the road, appearing along its entire 150 m length. The blocks were placed at three distinct distances relative to the centreline of the road, being either on the road (Road), near the road (Near), or further away (Far). An aerial view of the scene appears in Fig. 5(b). Although the blocks were placed randomly along the road, they were constrained to leave a clear path along which the driver could travel without collisions. Twelve subjects took part in the experiment and they were required to perform the detection task under all three locomotion conditions. The order in which they did each locomotion condition was permuted across subjects, such that exactly two subjects ran the experiment in any one of the six possible orders. In order to control the amount of the scene visible to both passive and active drivers, the path steered by the active drivers were recorded and then used to define the path along which a different subject moved as a passive observer. The choice of which subject received the active course from which other subject was random, although all trajectories were used exactly once.

Subject analysis of the scene was assessed by continuously changing the appearance of four of the sixty blocks in any one trial. There were four possible change types: Orientation, position, presence/absence, and colour. An orientation change resulted in a rotation of 10° about the vertical axis. A position change resulted in a shift in the ground plane of 0.5 m. A presence/absence change resulted in the complete disappearance of a block. A colour change produced a change in the block' s colour to one of three other colours (red, blue, yellow, and green).

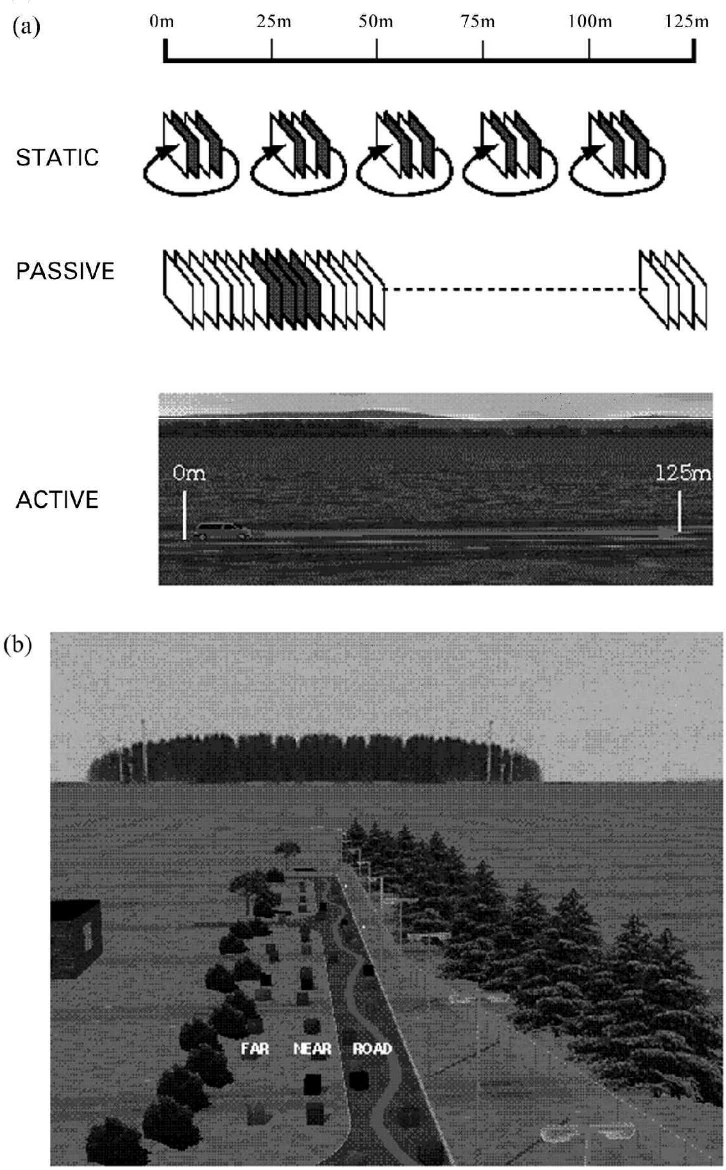

FIG. 5. (a) The three types of locomotion used in Experiment II. Subjects either saw the scene from five vantage points (Static), as a passively viewed smooth movement (Passive), or while steering themselves around obstacles in the road (Active). (b) An aerial view of the scene driven along by subjects. The curved line on the road represents the path driven by a particular subject.

The four changes were allowed to occur in any of the three locations. In other words all four changes might occur far from the road, or as few as none. The only constraint was that there were always a total of four changes per trial, and that the total number of times a certain number of changes occurred was the same for each of the three locations. Presenting four changes in this way resulted in 15 unique combinations of block location. The combinations were as follows: 4–0–0 (i.e. 4 on the road, none near to the road, and none far from the road), 3–0–1 (i.e. 3 on the road, none near to the road, and 1 far from the road), etc. Permuting all four change types with the 15 location combinations would have resulted in 15.(4!) = 360 trials. Since 360 trials would have taken over an hour to complete, we instead elected to run half this number, namely 180 trials. The choice of what type of change occurred was then chosen at random, although the total number of each change type was balanced over the trial block as a whole. This meant in practice that more than one colour, position, orientation, or presence change could occur in any one trial, but that the same number (180 of each) occurred over the complete set of trials. The 180 trials were carried out a total of three times, once under each type of the three locomotion conditions. During each pass along the road, subjects were asked to press a switch each time they noticed a block changing, and were then asked to describe the changes at the end of the trial. Apparent motion signals were subdued by obscuring the scene for 9 frames (120 msec at 75 Hz) with a uniform grey image. This occurred after 30 normal frames, and directly before each object change.

Results

A summary of the results appear in Fig. 6. Once again it is clear that subjects were unable to notice all of the changes which were made, and that their representation of the scene was lacking in sufficient detail to notice these changes even in the absence of ego-motion. Irrespective of locomotion condition, changes were often either missed or falsely categorized. A three-way within subjects ANOVA was used to analyse percentage correct with locomotion, change location, and change type as independent variables. All tests for independence of individual condition means were then made using Tukey' s Honestly Significant Difference Test.

All three factors showed significant effects—see Fig. 6. There was a clear effect of locomotion, $F(2,10) = 16.38$, $MSe = 487.20$, $p < .01$. Change location also affected detection significantly, $F(2,10) = 19.71$, $MSe = 130.26$, $p < .001$, as did change type, $F(3,15) = 18.47$, $MSe = 126.78$, $p < .001$. Average detection rates for the three conditions also dropped systematically with level of interaction. The number of changes seen per trial (N_C) under each condition was: Static $N_C = 3.07$, Passive $N_C = 2.5$, and Active $N_C = 2.15$.

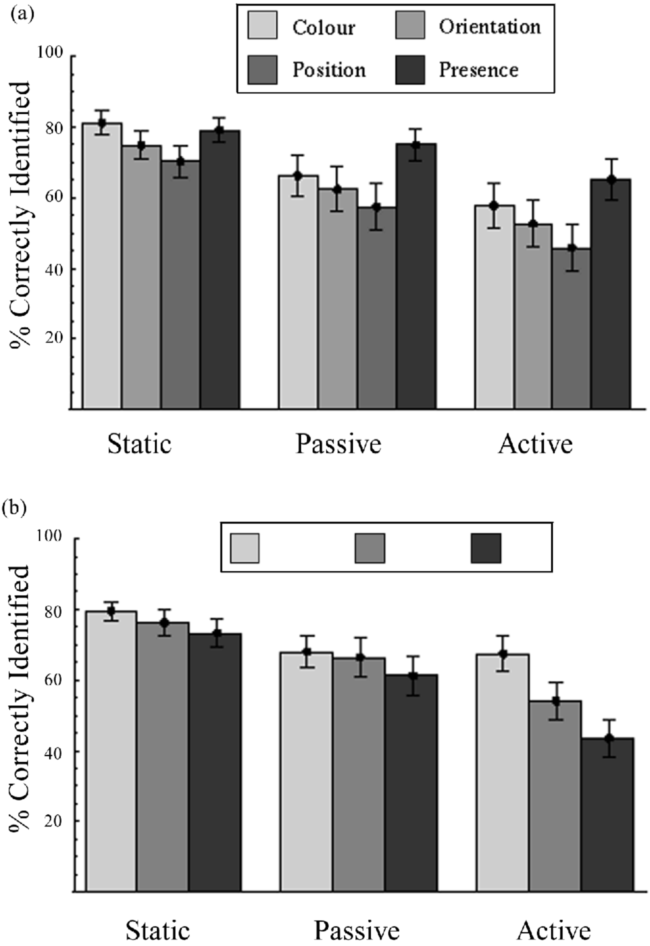

FIG. 6. The results of detecting changes to four of sixty blocks, separated according to whether the subject saw static images (Static), was moved through the scene (Passive), or steered themselves through the scene (Active). (a) Influence of the type of change seen. (b) Arranged by location of the change relative to the centre-line of the road. Error bars indicate *SE* of the mean in each case.

In practice, it is somewhat difficult to assess the main effects in this experiment, since there are several rather trivial explanations for why they occurred. First, the intrinsic detectability of changes was not controlled, that is, one should not be surprised if a change from red to yellow is not equivalent to a movement of 0.5 m or the complete disappearance of an object. Second, one might also not be surprised to see that the performance of subjects while steering was overall worse than in either of the other two conditions, because they are now partially occupied with avoiding obstacles. Third, it is to be expected that blocks on the road can be better analysed than those situated further away, due to perspective size reduction and occlusion.

Instead, it is more informative to consider how the various conditions interact. From Fig. 6(a), it is clear that the detection of orientation, position, and colour changes is influenced more by observer motion than presence changes, although the effect for colour is weaker than for the other two change types. Further analysis of the interaction terms revealed the associated effect. First, in the interaction between locomotion and change type, $F(6,30) = 3.33$, $MSe = 42.41$, $p < .05$, presence changes were significantly easier to detect than either position or orientation changes in both the Active and Passive conditions ($p < .05$). However, in the Static condition the difference was not significant, suggesting that position and orientation are more strongly influenced by observer motion than either colour or presence changes. There was also a significant interaction between locomotion and change location, $F(4,20) = 11.85$, $MSe = 47.27$, $p < .05$. All three location conditions differed significantly in the Active condition ($p < .05$), but not under the other forms of locomotion. This strongly suggests that the task of driving resulted in a narrowing of the main field of attention to the vicinity of the road, at the expense of other regions in the scene.

Conclusions

There are two important points to draw from the experiment. First, that the difference between the (Road, Near, Far) conditions, although visible under all modes of locomotion, only became significant when the subject steered. Performance when driving was generally worse than in the other two conditions, but due to a narrowing of attention to the road, change detection in that region was relatively unaffected compared to detection away from the road, which was much poorer. The second conclusion is that the difference between orientation and position changes to the other change types, only became significant when the subject moved, that is, position and orientation are less well represented when the subject is in motion.

DISCUSSION

Both experiments suggest that the impression we have of possessing a full and detailed mental representation of our surroundings is illusory. This adds strong support to the large body of evidence which has already reported similar effects in static scenes and movies, suggesting that we store remarkably little information about the position, colour, or size etc. of objects not currently under direct scrutiny. Results of this type have lead theorists such as O'Regan (1992), to suggest that the brain may simply continuously refixate certain key objects to enhance an otherwise poor representation held in memory. Precisely what would be included in this rough description is still unclear but it might retain some information about the identity and location of objects sufficient to build the scene' s context and to motivate later fixations—as described by several researchers (Hochberg, 1968; Irwin, Brown, & Sun, 1988; Ullman, 1984). The fact that the results presented here reveal particular sensitivity to the appearance or disappearance of objects certainly accords with this idea of an internalized object list (see also Mondy & Coltheart, this issue), but one in which the location of an object is only very roughly represented, possibly through its approximate relationship to others.

There are several developments in this study worthy of discussion. The first is the use of multiple changes. By introducing multiple changes, change detection tells the experimenter the relative salience of particular attributes and objects independent of the particular scene. Solving a change detection task actually often becomes a methodical serial search task, in which the observer fixates individual objects and waits for the state flip—indicated by the grey flash. This introduces a new, somewhat artificial context to the scene—namely the expectation that something will change. In their studies Hayhoe, Bensinger, and Ballard (1998) avoided this by not telling their subjects that something would change, and by giving them, as we did, a natural task to do. This obviously works only up to the point at which the "cat is out of the bag" and the subjects first become aware that things are changing. In our experiments, the subjects were told to look for changes, but by including many changes it was hoped that the effect of serial search could be reduced by comparing change detection rates within the same trial and same scene.

Another important development in these experiments is obviously the inclusion of ego motion. When in motion the colour or presence of an object does not change greatly from one instant to the next. In contrast, the object' s perceived orientation and location does, and the results described here reveal a commensurate decrease in detection performance for orientation and location changes. On the other hand, although subjects found it more difficult to detect view and position changes, they were able to predict how an object being studied should transform over a 200 msec interval. At least they were able to detect inconsistencies with that prediction when the size of the inconsistency was of the order used in both of the experiments described here.

One question that remains unanswered but which follows directly from this study of self-motion, is how well changes can be detected in a truly dynamic scene, in which objects within the scene can move too. Presumably this would make change detection even more difficult than in the self-motion case, although this has still to be tested. It would, for instance, be interesting to see if sudden changes to velocity or acceleration profiles of objects can also be detected. Our ability to avoid or catch moving targets, suggests that we are able to process these characteristics, but how many and with what sensitivity remains an open question.

Apart from the inclusion of motion in the scene, we also considered how task affects detection performance, and were able to show that by introducing an active steering task, subjects noticed changes to objects on the road more readily than off it. It seems likely that this is due both to attention and the choice of saccade targets in the scene, providing further proof that the attributes of unattended objects are poorly represented. It also suggests that objects may never be afforded detailed processing if they are irrelevant to the task in hand, as described by Hayhoe (this issue) as well. The fact that the change detection technique is sensitive to this switch in the subject' s field of interest underlines the belief of ourselves and others writers in this issue, that change detection is a powerful technique for assessing what people actually see when they look around them.

REFERENCES

Aginsky, V., & Tarr, M.J. (this issue). How are different properties of a scene encoded in visual memory? *Visual Cognition, 7,* 147–162.

Biederman, I., Mezzanotte, R.J., & Rabinowitz, J.C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, *14*, 143–177.

Blackmore, S.J., Brelstaff, G., Nelson, K., & Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, *24*, 1075–1081.

Boyce, S.J., & Pollatsek, A. (1992). Identification of objects in scenes: The role of scene background in object naming. *Journal of Experimental Psychology: LMC*, *18*, 531–543.

Boyce, S.J., Pollatsek, A., & Rayner, K. (1989). Effect of background information on object identification. *Journal of Experimental Psychology: HPP*, *15*, 556–566.

Carr, K., & England, R. (1993). *Simulated and virtual realities: Elements of perception*. London: Taylor & Francis.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (Ed.), *Vancouver studies in cognitive sciences perception: Vol. 5*. (pp. 89–110). Oxford, UK: Oxford University Press.

Hayhoe, M. (this issue). Vision using routines: A functional account of vision. *Visual Cognition, 7,* 43–64.

Hayhoe, M.M., Bensinger, D.G., & Ballard, D.H. (1998). Task constraints in visual working memory. *Vision Research*, *38*, 127–135.

Hochberg, J. (1968). In the mind' s eye. In R.N. Haber (Ed.), *Contemporary theory and research in visual perception* (pp. 309–331). New York: Holt, Reinhart & Winston.

Hochberg, J. (1986). Representation of motion and space in video and cinematic displays. In K.R. Boff, L.K. Kaufman, & J.P. Thomas (Eds), *Handbook of perception and human performance*, (pp. 22.1–22.64). London: John Wiley & Sons.

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

Irwin, D., Brown, J., & Sun, J. (1988). Visual masking and visual integration across saccadic eye movements. *Journal of Experimental Psychology, General*, *117*, 276–287.

Kalawsky, R.S. (1993). *The science of virtual reality and virtual environments*. Reading, MA: Addison-Wesley.

Larish, J.F., & Andersen, G.J. (1995). Active control in interrupted dynamic spatial orientation: The detection of orientation change. *Perception and Psychophysics*, *57*, 533–545.

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

McConkie, G.W., & Currie, C.B. (1996). Visual stability across saccades while viewing complex pictures. *Journal of Experimental Psychology: HPP*, *22*, 563–581.

McConkie, G.W., & Zola, D. (1979). Is visual information integrated across successive fixations in reading? *Perception and Psychophysics*, *25*, 221–224.

Mondy, S., & Coltheart, V. (this issue). Detection and identification of change in naturalistic scenes. *Visual Cognition, 7,* 281–296.

Mourant, R.R., & Rockwell, T.H. (1972). Strategies of visual search by novice and experienced drivers. *Human Factors*, *14*, 325–335.

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., Deubel, H., Clark, J.J., & Rensink, R.A. (this issue). Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition, 7,* 191–211.

Péruch, P., Vercher, J.L., & Gauthier, G.M. (1995). Acquisition of spatial knowledge through visual exploration of simulated environments. *Ecological Psychology*, *7*, 1–20.

Previc, F.H. (1990). Functional specialisation in the upper and lower fields in humans: Its ecological origins and neurophysiological implications. *Behavioural and Brain Sciences*, *13*, 519–575.

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

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., & Levin, D.T. (1997). Failure to detect changes to attended objects in motion pictures. *Trends in Cognitive Sciences*, *1*, 261–267.

Telford, L., Howard, I.P., & Ohimi, M. (1995). Heading judgements during active and passive self-motion. *Experimental Brain Research*, *104*, 502–510.

Ullman, S. (1984). Visual routines. *Cognition*, *18*, 97–159.

Wilson, P.N., Foreman, N., Gillett, R., & Stanton, D. (1997). Active versus passive processing of spatial information in a computer-simulated environment. *Ecological Psychology*, *9*, 207–222.