

Gorillas in our midst: sustained inattentional blindness for dynamic events

Daniel J Simons, Christopher F Chabris

Department of Psychology, Harvard University, 33 Kirkland Street, Cambridge, MA 02138, USA;

e-mail: dsimons@wjh.harvard.edu

Received 9 May 1999, in revised form 20 June 1999

Abstract. With each eye fixation, we experience a richly detailed visual world. Yet recent work on visual integration and change direction reveals that we are surprisingly unaware of the details of our environment from one view to the next: we often do not detect large changes to objects and scenes ('change blindness'). Furthermore, without attention, we may not even perceive objects ('inattentional blindness'). Taken together, these findings suggest that we perceive and remember only those objects and details that receive focused attention. In this paper, we briefly review and discuss evidence for these cognitive forms of 'blindness'. We then present a new study that builds on classic studies of divided visual attention to examine inattentional blindness for complex objects and events in dynamic scenes. Our results suggest that the likelihood of noticing an unexpected object depends on the similarity of that object to other objects in the display and on how difficult the priming monitoring task is. Interestingly, spatial proximity of the critical unattended object to attended locations does not appear to affect detection, suggesting that observers attend to objects and events, not spatial positions. We discuss the implications of these results for visual representations and awareness of our visual environment.

1 Introduction

"It is a well-known phenomenon that we do not notice anything happening in our surroundings while being absorbed in the inspection of something; focusing our attention on a certain object may happen to such an extent that we cannot perceive other objects placed in the peripheral parts of our visual field, although the light rays they emit arrive completely at the visual sphere of the cerebral cortex."

Rezső Bálint 1907 (translated in Husain and Stein 1988, page 91)

Perhaps you have had the following experience: you are searching for an open seat in a crowded movie theater. After scanning for several minutes, you eventually spot one and sit down. The next day, your friends ask why you ignored them at the theater. They were waving at you, and you looked right at them but did not see them. Just as we sometimes overlook our friends in a crowded room, we occasionally fail to notice changes to the appearance of those around us. We have all had the embarrassing experience of failing to notice when a friend or colleague shaves off a beard, gets a haircut, or starts wearing contact lenses. We feel that we perceive and remember everything around us, and we take the occasional blindness to visual details to be an unusual exception. The richness of our visual *experience* leads us to believe that our visual *representations* will include and preserve the same amount of detail (Levin et al 2000).

The disparity between the richness of our experience and the details of our representation, though 'well known' to Bálint in 1907, has been studied only sporadically in the psychological literature since then, and many of the most striking results appear to have been neglected by contemporary researchers. Although the past 20 years have seen increasing interest in the issue of the precision of visual representations, a series of studies from the 1970s and 1980s using dynamic visual displays provides some of the most dramatic demonstrations of the importance of attention in perception (see Neisser 1979 for an overview). In these studies, observers engage in a continuous task that requires them to focus on one aspect of a dynamic scene while ignoring others.

At some point during the task an unexpected event occurs, but the majority of observers do not report seeing it even though it is clearly visible to observers not engaged in the concurrent task (Becklen and Cervone 1983; Littman and Becklen 1976; Neisser 1979; Neisser and Becklen 1975; Rooney et al 1981; Stoffregen et al 1993; Stoffregen and Becklen 1989). Although these studies have profound implications for our understanding of perception with and without attention, and despite their obvious connection to more recent work on visual attention (eg change blindness, attentional blink, repetition blindness, inattention blindness), the empirical approach has fallen into disuse. One goal of our research is to revive the approach used in these original studies of 'selective looking' in the context of more recent work on visual attention.

Over the past few years, several researchers have demonstrated that conscious perception seems to require attention. When attention is diverted to another object or task, observers often fail to perceive an unexpected object, even if it appears at fixation—a phenomenon termed 'inattention blindness' (eg Mack and Rock 1998).⁽¹⁾ These findings are reminiscent of another set of findings falling under the rubric of 'change blindness'. Observers often fail to notice large changes to objects or scenes from one view to the next, particularly if those objects are not the center of interest in the scene (Rensink et al 1997). For example, observers often do not notice when two people in a photograph exchange heads, provided that the change occurs during an eye movement (Grimes 1996; see Simons and Levin 1997 for a review). Such studies suggest that attention is necessary for change detection (see also Scholl 2000), but not sufficient, as even changes to attended objects are often not noticed (Levin and Simons 1997; Simons and Levin 1997, 1998; Williams and Simons 2000). For example, observers who were giving directions to an experimenter often did not notice that the experimenter was replaced by a different person during an interruption caused by a door being carried between them (Simons and Levin 1998).

Both areas of research focus on two fundamental questions. (i) To what degree are the details of our visual world perceived and represented? (ii) What role does attention play in this process? We will review recent evidence for inattention blindness to provide a current context for a discussion of earlier research on the perception of unexpected events. We then present a new study examining the variables that affect inattention blindness in naturalistic, dynamic events, and consider the results within the broader framework or recent attention research, including change blindness.

1.1 *Inattention blindness*

Studies of change blindness assume that, with attention, features can be encoded (abstractly or otherwise) and retained in memory. That is, all of the information in the visual environment is potentially available for attentive processing. Yet, without attention, not much of this information is retained across views. Studies of inattention blindness have made an even stronger claim: that, without attention, visual features of our environment are not perceived at all (or at least not consciously perceived)—observers may fail not just at change detection, but at perception as well.

Recent work on the role of attention in perception has explored what happens to unattended parts of simple visual displays (Mack and Rock 1998; Mack et al 1992; Moore and Egeth 1997; Newby and Rock 1998; Rock et al 1992; Rubin and Hua 1998; Silverman and Mack 1997). In traditional models of visual search, features are often assumed to be processed preattentively if search speeds are unaffected by the number

⁽¹⁾ Mack and Rock (1998) draw a distinction between conscious perception and implicit perception. Consistently with this distinction, when we use the term 'perceive' (or 'notice' or 'see') in this paper, we mean that observers have at some point had a conscious experience of an object or event. However, it is important to note that even when observers do not perceive an object, it may still have an implicit influence on their subsequent decisions and performance (eg Chun and Jiang 1998; Moore and Egeth 1997).

of distracter items in the display (ie the feature ‘pops out’ effortlessly). Preattentive processing of some features would allow for rapid perception of more complex objects that are built by combining such sensory primitives. However, visual search tasks may not truly assess the processing of unattended stimuli because observers have the expectation that a target may appear—observers know that they will have to search the display for a particular stimulus. Hence, they may *expect* to perceive these features, which would allow their visual/cognitive system to anticipate the features. The inattentional-blindness paradigm developed by Mack, Rock, and colleagues avoids this potential confound of knowledge of the task (eg Mack and Rock 1998), allowing a more direct assessment of the perception of unattended stimuli. In a typical version of their task, observers judge which of two arms of a briefly displayed large cross is longer. On the fourth trial of this task, an unexpected object appears at the same time as the cross. After this trial, observers are asked to report if they saw anything other than the cross. After answering this question, observers view another trial, now with the suggestion that something might appear. This allows an assessment of perception under conditions of divided attention. Last, subjects complete a final, full-attention trial in which they look for and report the critical object but ignore the cross. Performance on the critical, unattended trial is compared with that on the divided-attention and full-attention trials to estimate the degree to which attention influences perception. The difference in the proportion of subjects noticing on the full-attention and critical trials is the amount of inattentional blindness.

Several clear patterns emerge from this body of research (see Mack and Rock 1998 for an overview). (i) About 25% of subjects are inattentionally blind when the cross is presented at fixation and the unexpected object is presented parafoveally (subjects typically detect the critical stimulus on divided-attention and full-attention trials). (ii) About 75% of subjects are inattentionally blind when the cross is presented parafoveally and the unexpected object is presented at fixation, suggesting an effortful shift of attention away from fixation to the cross and possible inhibition of processing at the ignored fixation location. (iii) These levels of detection are no different for features thought to be preattentively processed (eg color, orientation, motion) and those thought to require effort. (iv) Although objects composed of simple visual features are not easily detected, some meaningful stimuli are. Observers typically notice their own name or a smiley face even when they did not expect it. Note, however, they do not tend to notice their own name if one letter is changed (see also Rubin and Hua 1998). Observers do not consciously perceive the visual features, but they do perceive the meaning. (v) Observers seem to focus attention on particular locations on the screen. Objects that appear inside this zone of attention are more likely to be detected than those appearing outside (Mack and Rock 1998; Newby and Rock 1998), suggesting that attention is focused not on the object or event itself, but on the area around that object.

1.2 ‘Selective looking’

These recent studies of inattentional blindness used simple, brief visual displays under precisely controlled timing conditions, in the vein of work on visual search and related attention paradigms that were largely designed to examine how we select and process features and objects. The paradigm was designed to be a visual analogue of dichotic-listening studies conducted during the 1950s and 1960s (Cherry 1953; Moray 1959; Treisman 1964), and largely succeeded in replicating the classic auditory effects with visual stimuli. Although relatively little unattended information reaches awareness, some particularly meaningful stimuli do. Despite the similarity of these theoretical conclusions, they are fundamentally different in an important way. Almost by necessity, dichotic-listening tasks involve dynamic rather than static events. Listening studies reveal a degree of ‘inattentional deafness’ that extends over time and over changes in the unattended stimulus.

In that sense, the computer-based inattention paradigm is not a true analogue of dichotic-listening tasks. Although the theoretical conclusions match our experience of not seeing friends in a crowded theater (and hearing our own name spoken at a noisy party), the experimental paradigm may not fully capture all aspects of that natural situation [see Neumann et al (1986) for a discussion of the difficulties of equating auditory and visual divided-attention tasks]. However, an earlier series of studies by Neisser and his colleagues did use dynamic events to address many of the same questions.

In an initial study (Neisser and Becklen 1975), observers viewed a display which presented two overlapping, simultaneous events. (The superimposition was achieved by showing both of the separately recorded events on an angled, half-silvered mirror.) One of the events was a hand-slapping game in which one player extended his hands with palms up and the other player placed his hands on his opponents hands with palms down. The player with palms up tries to slap the back of the other player's hands, and the other player tries to avoid the slap. The second event depicted three people moving in irregular patterns and passing a basketball. Subjects were asked to closely monitor one of the two events. If they monitored the hand game, they pressed a button with each attempted slap. If they monitored the ball game, they pressed the button for each pass. Each subject viewed a total of ten trials. The first two trials showed each of the games alone. On the 3rd and 4th trials, both events were presented simultaneously, but subjects were asked to follow only one of them. On the 5th and 6th trials, subjects attempted to respond to both events, using one hand to respond to each (only twenty actions per minute rather than forty occurred in these two and subsequent trials). On the last four trials, subjects responded to only one of the events, but an additional unexpected event occurred as well. In trial 7, the two hand-game players stopped and shook hands. On trial 8, one of the ball-game players threw the ball out of the game and the players continued to pretend to be passing the ball. The ball was returned after 20 s of fake throws. On trial 9, the hand-game players briefly stopped their game and passed a small ball back and forth. On trial 10, each of the ball-game players stepped off camera and was replaced by a woman and, after 20 s, the original men returned in the same fashion.

The results of this study are largely consistent with the findings of computer-based inattention studies. In the initial trials, subjects could easily follow one event while ignoring another event occupying the same spatial position. [This was true even when subjects were not allowed to move their eyes; see Littman and Becklen (1976).] Not surprisingly, they had much greater difficulty simultaneously monitoring both events. More importantly, in the initial trial with an unexpected event, only one of twenty-four people spontaneously reported the hand shake, and three others mentioned it in post-experiment questioning. None of the subjects spontaneously reported the disappearance of the ball, three spontaneously reported the ball pass in the hand game, and three reported the exchange of women for men on the final trial. Subjects who noticed one of the unusual events were more likely to notice subsequent unusual events, much as subjects in the divided-attention conditions in inattentional-blindness studies typically reported the presence of the previously 'unexpected' object (Mack and Rock 1998). In total, 50% of Neisser and Becklen's (1975) subjects showed no indication of having seen any of the unexpected events, and even subjects who did notice typically could not accurately report the details of them.

In a more recent version of this sort of divided-visual-attention task, observers viewed superimposed videotapes of two of the ball games described above (Becklen, Neisser, and Littman, discussed in Neisser 1979).⁽²⁾ The players in one game wore

⁽²⁾ Many of the 'selective-looking' studies conducted by Neisser and his colleagues were never published in complete empirical reports. In such cases, as here, we have cited unpublished or in-preparation manuscripts on the basis of their descriptions in other, published materials.

black shirts and the players in the other game wore white shirts. This change made the attended and ignored events more similar, and therefore more difficult to discriminate. Nevertheless, observers could successfully follow one game while ignoring the other even when both teams wore the same clothing (in fact, the same three players appeared in each video stream).

In subsequent studies of selective looking, Neisser and his colleagues used this 'basketball-game' task [see Neisser (1979) for a description of several different versions]. In the most famous demonstration, observers attend to one team of players, pressing a key whenever one of them makes a pass, while ignoring the actions of the other team. After about 30 s, a woman carrying an open umbrella walks across the screen (this video was also superimposed on the others so all three events were partially transparent; see figure 1). She is visible for approximately 4 s before walking off the far end of the screen. The games then continue for another 25 s before the tape is stopped. Of twenty-eight naive observers, only six reported the presence of the umbrella woman, even when questioned directly after the task (Neisser and Dube, cited in Neisser 1979). Interestingly, when subjects had practice performing the task on two similar trials before the trial with the unexpected umbrella woman, 48% noticed her. When subjects just watched the screen and did not perform any task, they always noticed the umbrella woman, a result consistent with the inattentional-blindness findings reviewed earlier (and with work on saccade-contingent changes; see Grimes 1996; McConkie and Zola 1979).



Figure 1. A single frame captured from a late-generation video of the umbrella-woman sequence used by Neisser and colleagues (eg Neisser 1979). The woman is in the center of the image and her umbrella is white.

Interestingly, Neisser (1979) mentioned an additional study in which the umbrella woman wore the same-color shirt as either the attended or the unattended team. Apparently, this feature-similarity manipulation caused little difference in the rate of noticing. Also, when the unexpected character was a small boy rather than the umbrella woman, fewer subjects noticed him, and when the umbrella woman stopped her motion and performed a little dance, more subjects noticed. These latter two findings suggest that properties of an unattended stimulus can capture attention, even though similarity to attended stimuli seemed to make little difference. However, these findings must be evaluated tentatively, because the details of the experimental paradigm were not presented by Neisser (1979).

In one of the few published empirical reports with this paradigm, Becklen and Cervone (1983) examined the effect of eliminating the delay between the umbrella-woman's appearance and the questioning of subjects. They found no difference in noticing rates when the entire video was shown (35%) and when the video ended immediately after the umbrella woman left the scene (33%). Furthermore, performance was substantially worse when the video ended as the umbrella woman was halfway across the court (7%), even though that meant that the last image subjects saw included the umbrella woman. These subjects provided accurate descriptions of the scene, including details of the player locations, but did not mention an umbrella woman. Neisser and Rooney (cited in Becklen and Cervone 1983) addressed the same question by interrupting the action as the umbrella woman was two-thirds of the way across the screen. Immediately after the interruption, a split screen appeared with the umbrella woman on one side and a boy holding a soda can on the other; subjects were asked to pick which they had seen. When the results were corrected for guessing, only 30% of observers had noticed the umbrella woman, a level comparable to that shown in postexperiment interviews.

These findings provide important evidence against the notion of 'inattentional amnesia', an alternative account of findings of inattentional blindness and change blindness. According to this view (Wolfe 1999), the unexpected event is consciously perceived, but immediately forgotten. Hence, the failure to report its appearance reflects a failure of memory rather than of perception. In this case, however, even though subjects are tested immediately after the event, they are no better at detecting it. Furthermore, when people notice the unexpected event in this task, they sometimes smile or laugh; nonnoticers show no outward signs of detection. The forgetting would have to be so rapid as to be inseparable from the act of perception to allow any sort of amnesia to explain these findings.

This early work on selective looking raised a number of questions needing further study. What role does similarity between the unexpected and attended events play in detection? Are particularly unusual events more or less likely to be detected? Does task difficulty increase or decrease detection? Perhaps the most important question left unanswered in this early work is what role the unusual superimposition of the events played in causing inattentional blindness. Most cognitive psychologists we have talked to found these results interesting, but were somewhat less convinced of the importance of the failures to notice unexpected events. After all, the video superimposition gives the displays an odd appearance, one not typically experienced in the real world and one in which the players and the umbrella woman are not as easy to see as they would be without superimposition.

One more recent study has looked at performance when all of the actors and the umbrella woman are shot from a single video camera, with no superimposition (Stoffregen et al 1993). Under these conditions, the players and umbrella woman occluded each other and the balls. If failures to notice the umbrella woman in earlier studies resulted from the unnatural appearance of the superimposed version of the display, performance might be much better with a 'live' version. Subjects performed the task for approximately 30 s before the umbrella woman appeared and walked across the screen. The camera angle used for this film was much wider than in earlier studies—it showed an entire regulation basketball court. Consequently, the umbrella woman was visible for a longer time (12 s) and the players and the umbrella woman were substantially smaller on screen than in earlier studies. Another notable difference is that only twelve passes occurred during the 60-s video (rather than 20–40 as in earlier studies). Even in this live version of the study, only three of twenty normal subjects tested reported the presence of the umbrella woman. Although this finding does suggest that visual superimposition was not the cause of failures of noticing, it did

not match the stimulus conditions of the other studies and did not directly compare performance with and without superimposition. The difference in camera angle (and consequent character size) alone may well have affected detection rates, so this study is not a well-controlled test of the generalizability of inattentional-blindness phenomena to more natural stimulus conditions.

Despite the importance of all the unanswered questions raised by these studies, to our knowledge the findings reviewed above are the only published reports using dynamic, naturalistic events to study the detection of unexpected objects.⁽³⁾ Taken together, these studies lead to a number of striking conclusions, some consistent and others inconsistent with findings with simple displays. Unlike the computer-based studies (eg Mack and Rock 1998; Newby and Rock 1998), the video studies demonstrate that inattentional blindness does not result from attention being focused *elsewhere* in the display. In the superimposed version of the display, the umbrella woman occupied exactly the same spatial position as the attended players and balls. In fact, the balls even passed through the umbrella woman. This finding is inconsistent with the computer-based result that detection was better when the unexpected object appeared within the region defined by the attended object (Mack and Rock 1998). Several factors might account for this difference. First, there were simply more objects to attend to in the video displays, so attention may not have stayed on any one location for long. Second, the dynamic display may have captured and held attention more effectively than the cross task. Third, the video task may simply have been harder, leaving fewer attentional resources available to process unanticipated events. These video studies do show that a form of inattentional blindness can last much longer than the brief exposure times used in recent static-display studies. Subjects missed ongoing events that lasted for more than 4 s.

Although these differences between the computer-based and video studies are important, the general similarity of the conclusions is striking. In both cases, observers often do not see unanticipated objects and events. The video studies suggest that these findings can help explain real-world phenomena such as our inability to see our friends in a crowded movie theater or airplanes on an approaching runway when our attention is focused on a different goal. Both change blindness and inattentional blindness show that attention plays a critical role in perception and in representation. Without attention, we often do not see unanticipated events, and even with attention, we cannot encode and retain all the details of what we see.

Although these video studies of inattentional blindness help to generalize findings from simple displays to more complex situations, the original reports do not fully examine all of the critical questions. For example, there is a hint that the visual similarity of the unexpected object to the attended ones makes no difference, but the details of that study were never published. Furthermore, the experiments did not systematically consider the role of task difficulty in detection. Perhaps most importantly, no direct comparisons were made between performance with the superimposed version of the display and with the 'live' version. In the studies reported here, we attempt to examine each of these factors. We also consider the nature of the unusual event. To combine all of these factors orthogonally within a single consistent paradigm, we filmed several video segments with the same set of actors in the same location on the same day. We then asked a large number of naive observers to watch the video recordings and later answer questions about the unexpected events.

⁽³⁾ Haines (1989) did address this topic as part of a larger human-interface study. Pilots attempted to land a plane in a flight simulator while using a head-up display of critical flight information superimposed on the 'windshield'. Under these conditions, some pilots failed to notice that a plane on the ground was blocking their path. In addition, Mack and Rock (1998) report several studies in which the unexpected object moved stroboscopically across part of the display, often without being detected during the 200 ms viewing period.

2 Method

2.1 Observers

228 observers, almost all undergraduate students, participated in the experiment. Each observer either volunteered to participate without compensation, received a large candy bar for participating, or was paid a single fee for participating in a larger testing session including another, unrelated experiment.

2.2 Materials


Four videotapes, each 75 s in duration, were created. Each tape showed two teams of three players, one team wearing white shirts and the other wearing black shirts, who moved around in a relatively random fashion in an open area (approximately 3 m deep \times 5.2 m wide) in front of a bank of three elevator doors. The members of each team passed a standard orange basketball to one another in a regular order: player 1 would pass to player 2, who would pass to player 3, who would pass to player 1, and so on. The passes were either bounce passes or aerial passes; players would also dribble the ball, wave their arms, and make other movements consistent with their overall pattern of action, only incidentally looking directly at the camera.

After 44–48 s of this action, either of two unexpected events occurred: in the Umbrella-Woman condition, a tall woman holding an open umbrella walked from off camera on one side of the action to the other, left to right. The actions of the players, and this unexpected event, were designed to mimic the stimuli used by Neisser and colleagues. In the Gorilla condition, a shorter woman wearing a gorilla costume that fully covered her body walked through the action in the same way. In either case, the unexpected event lasted 5 s, and the players continued their actions during and after the event.

There were two styles of video: in the Transparent condition, the white team, black team, and unexpected event were all filmed separately, and the three video streams were rendered partially transparent and then superimposed by using digital video-editing software. (Neisser and colleagues achieved similar effects using analog equipment or a physical apparatus that superimposed separate displays by means of mirrors.) In the Opaque condition, all seven actors were filmed simultaneously and could thus occlude one another and the basketballs; this required some rehearsal to eliminate collisions and other accidents, and to achieve natural-looking patterns of movement. All videos were filmed with an SVHS video camera (Panasonic AG456U) and were digitized and edited by using a nonlinear digital-editing system (Media 100LX and Adobe Aftereffects, running on Power Computing hardware). All editing of the videos was accomplished after digitization, so the degree of signal loss due to multiple generations of editing was minimized and also equated across conditions. Stimuli were created by mastering the digitally edited sequences to VHS format tapes. Thus, as shown in figure 2, videos were created with the following four display types: Transparent/Umbrella Woman, Transparent/Gorilla, Opaque/Umbrella Woman, and Opaque/Gorilla. The first of these was most similar to the conditions tested by Neisser and colleagues.

2.3 Procedures

All observers were tested individually and gave informed consent in advance. Before viewing the videotape, observers were told that they would be watching two teams of three players passing basketballs and that they should pay attention to either the team in white (the White condition) or the team in black (the Black condition). They were told that they should keep either a silent mental count of the total number of passes made by the attended team (the Easy condition) or separate silent mental counts of the number of bounce passes and aerial passes made by the attended team (the Hard condition). Thus, for each of the four displays, there were four task conditions—White/Easy, White/Hard, Black/Easy, and Black/Hard—for a total of sixteen individual conditions. Each observer participated in only one condition.



Transparent/Umbrella Woman



Transparent/Gorilla



Opaque/Umbrella Woman



Opaque/Gorilla

Figure 2. Single frames from each of the display tapes used here. (These tapes and that referred to in figure 3 were in color. These frames are displayed in color on <http://www.perceptionweb/perc0999/simons.htm/> and archived on the CD ROM supplied with issue 12 of *Perception*.) The transparent conditions (top row) were created by superimposing three separately filmed events by means of digital video editing. The opaque conditions (bottom row) were filmed as a single action sequence with all seven actors. This figure shows the display for each condition halfway through the unexpected event, which lasted for 5 s of the 75-s-long video.

After viewing the videotape and performing the monitoring task, observers were immediately asked to write down their count(s) on paper.⁽⁴⁾ They were then asked to provide answers to a surprise series of additional questions. (i) While you were doing the counting, did you notice anything unusual on the video? (ii) Did you notice anything other than the six players? (iii) Did you see anyone else (besides the six players) appear on the video? (iv) Did you see a gorilla [woman carrying an umbrella] walk across the screen? After any “yes” response, observers were asked to provide details of what they noticed. If at any point an observer mentioned the unexpected event, the remaining questions were skipped. After the questioning, observers were asked whether they had ever previously participated in an experiment similar to this or had ever heard of such an experiment or the general phenomenon. (Observers who answered “yes” were replaced and their data were discarded.) Last, the observer was debriefed; this included replaying the videotape on request. Each testing session lasted 5–10 min.

Twenty-one experimenters tested the observers. To ensure uniformity of procedures, we developed a written protocol in advance and reviewed it with the experimenters before they began to collect data. This document specified what the experimenters would say to

⁽⁴⁾Note that in all the Transparent conditions, the correct counts were identical because the same passing sequences were used to create both of the Transparent display tapes (Umbrella Woman and Gorilla). In the Opaque conditions, the correct counts varied because the passing sequences were filmed separately for each of the unexpected events.

each observer, when they would say it, how and when they would show the videotape, how they would collect and record the data, and how they would debrief observers. Experimenters used a variety of television monitors, ranging from 13 to 36 inches (diagonal) in screen size to present the videotapes.

3 Results

Data from thirty-six observers were discarded for a variety of reasons: either (i) the observer already knew about the phenomenon and/or experimental paradigm ($n = 14$), (ii) the observer reported losing count of the passes ($n = 9$), (iii) passes were incompletely or inaccurately recorded ($n = 7$), (iv) the observer's answer could not be clearly interpreted ($n = 5$), or (v) the observer's total pass count was more than three standard deviations away from the mean of the other observers in that condition ($n = 1$). The remaining 192 observers were distributed equally across the sixteen conditions of the $2 \times 2 \times 2 \times 2$ design (twelve per condition).

Although we asked a series of questions escalating in specificity to determine whether observers had noticed the unexpected event, only one observer who failed to report the event in response to the first question ("did you notice anything unusual?") reported the event in response to *any* of the next three questions (which culminated in "did you see a ... walk across the screen?"). Thus, since the responses were nearly always consistent across all four questions, we will present the results in terms of overall rates of noticing. Table 1 shows these results for each of the sixteen conditions.

Table 1. Percentage of subjects noticing the unexpected event in each condition. Each row corresponds to one of the four video display types. Columns are grouped by monitoring task and attended team (White or Black). In the Easy task, subjects counted the total number of passes made by the attended team. In the Hard task, subjects maintained separate simultaneous counts of the aerial and bounce passes made by the attended team.

	Easy task		Hard task	
	White team	Black team	White team	Black team
Transparent				
Umbrella Woman	58	92	33	42
Gorilla	8	67	8	25
Opaque				
Umbrella Woman	100	58	83	58
Gorilla	42	83	50	58

Out of all 192 observers across all conditions, 54% noticed the unexpected event and 46% failed to notice the unexpected event, revealing a substantial level of sustained inattentive blindness for a dynamic event and confirming the basic results of Neisser and colleagues. More observers noticed the unexpected event in the Opaque condition (67%) than in the Transparent condition (42%); $\chi^2_1 = 12.084$, $p < 0.001$; $n = 96$ per condition. However, even in the Opaque case, a substantial proportion of observers (33%) failed to report the event, despite its visibility and the repeated questions about it.

More observers noticed the unexpected event in the Easy (64%) than in the Hard (45%) condition ($\chi^2_1 = 6.797$, $p < 0.009$; $n = 96$ per condition). To confirm that these monitoring tasks differed in difficulty, we calculated the SD of the total pass counts reported by observers in each condition; the average SD was 2.71 in the eight Easy conditions and 6.77 in the eight Hard conditions, indicating that the Hard monitoring task was indeed more difficult. Accordingly, the correlation across conditions between the frequency of noticing (shown in table 1) and the SD of the total pass count was $r = -0.56$. The effect of task difficulty was greater in the Transparent conditions (Easy 56%, Hard 27%; $\chi^2_1 = 8.400$, $p < 0.004$; $n = 48$ per condition) than in

the Opaque conditions (Easy 71%, Hard 62%; $\chi^2_1 = 0.750$, $p < 0.386$; $n = 48$ per condition), suggesting a multiplicative effect on residual attention capacity of tracking difficult-to-see stimuli and keeping two running counts in working memory.

Next we examined differences in the detection of the two unexpected events. The Umbrella Woman was noticed more often than the Gorilla overall (65% versus 44%; $\chi^2_1 = 8.392$, $p < 0.004$; $n = 96$ per condition). This relation held regardless of the video type, monitoring task, or attended team, suggesting that the Umbrella Woman was either a more visually salient event than the Gorilla,⁽⁵⁾ more consistent with observers' expectations about situations involving basketballs, more semantically similar to the attended events, or all three. However, when observers attended to the actions of the Black team, they noticed the Gorilla much more often than when they attended to the actions of the White team (Black 58%, White 27%; $\chi^2_1 = 9.579$, $p < 0.002$; $n = 48$ per condition). By contrast, attending the Black team versus the White team made little difference in noticing the Umbrella Woman (Black 62%, White 69%; $\chi^2_1 = 0.416$, $p < 0.519$; $n = 48$ per condition). The Gorilla was black, whereas the Umbrella Woman wore pale colors that differed from both the Black and the White team. Thus, contrary to the suggestion of Neisser (1979), it appears that observers are more likely to notice an unexpected event that *shares* basic visual features—in this case, color—with the events they are attending to. In a sense, this effect is the opposite of the traditional 'pop-out' phenomenon in visual search tasks, which occurs when an item that *differs* in basic visual features from the rest of the display is easier to notice and identify.

It is possible that subjects who lost count of the passes would be most likely to notice the unexpected event. This is unlikely, however, for two reasons. First, subjects who reported losing count were replaced prior to data analysis. Second, we calculated the point-biserial correlation r between noticing (coded as 1 for reporting and 0 for not reporting the event) and the subject's absolute deviation from an accurate pass count (measured as the number of passes above or below the correct range) for each condition except the Opaque/Umbrella-Woman/White/Easy condition, which engendered 100% noticing. Across these fifteen conditions the correlations averaged to $r = 0.15$, suggesting that noticing was not strongly associated with counting poorly or inattentively.

4 Discussion

Our findings have replicated, generalized, and extended the surprising result first reported by Neisser and colleagues (Bahrick et al 1981; Becklen and Cervone 1983; Littman and Becklen 1976; Neisser 1979; Neisser and Becklen 1975; Rooney et al 1981; Stoffregen et al 1993; Stoffregen and Becklen 1989), and have demonstrated a robust phenomenon of sustained inattentive blindness for dynamic events. In particular, we have shown the following.

(i) Approximately half of observers fail to notice an ongoing and highly salient but unexpected event while they are engaged in a primary monitoring task. This extends the phenomenon of inattentive blindness (eg Mack and Rock 1998) by at least an order of magnitude in the duration of the event that can be missed. To stretch this limit further, we tested a longer and more salient unexpected event in an additional condition not reported above. In a separate Opaque-style video recording, the Gorilla walked from right to left into the live basketball-passing event, stopped in the middle of the players as the action continued all around it, turned to face the camera, thumped its chest, and then resumed walking across the screen (this action began after 35 s and lasted 9 s in a stimulus tape 62 s long; see figure 3 for a still frame).

⁽⁵⁾Visual salience, here, could refer to the relative distinctiveness of the unexpected objects in relation to the other players or to the background of the scene. Furthermore, the Umbrella Woman may have been spatially more distinctive in that her umbrella extended above the heads of the other players whereas the Gorilla was the same height as the other players.




Figure 3. A single frame from an additional experimental condition in which the gorilla stopped in the middle of the display, turned to face the camera, thumped its chest, and then continued walking across the field of view. Subjects performed the Easy monitoring task while attending to the White team, and the noticing rate was similar to that in the corresponding condition with the standard (shorter) Opaque/Gorilla event.

Twelve new observers⁽⁶⁾ watched this video while attending to the White team and engaging in the Easy monitoring task; only 50% noticed the event. This is roughly the same as the percentage that noticed the normal Opaque/Gorilla-walking event (42%) under the same task conditions.

(ii) This sustained inattentional blindness occurs more frequently if the display is transparent, with actors seeming to move through each other (as used in earlier studies), but observers often miss even fully visible objects appearing in live-action opaque displays. This latter finding is contrary to the intuitions of researchers who believed that the original effect was due to the unusual nature of the transparent video, and provides further evidence that inattentional blindness is a ubiquitous perceptual phenomenon rather than an artifact of particular display conditions.

(iii) The level of inattentional blindness depends on the difficulty of the primary task; in principle, inattentional blindness in this paradigm could be continuously varied by appropriately manipulating the difficulty of the monitoring task.

(iv) Observers are more likely to notice unexpected events if these events are visually similar to the events they are paying attention to. (On the basis of our results it is logically possible that dissimilarity to the ignored events is instead the crucial factor.)

(v) Objects can pass through the spatial extent of attentional focus (and the fovea) and still not be 'seen' if they are not specifically being attended. This conclusion is consistent with Mack and Rock's (1998) finding that observers often fail to notice a bar or square moving stroboscopically across fixation during a 200 ms display. In each of our videotapes, the unexpected object more than once crossed the path of the basketball and/or a player throwing or catching the ball, the observers would have had to pay attention to both of those elements of the display to perform the monitoring task.

In most respects, the results of this study are consistent with computer-based studies of inattentional blindness. Observers fail to report unexpected, suprathreshold objects when they are engaged in another task. Both sets of findings are consistent with the

⁽⁶⁾Data from two additional observers were discarded, one because he already knew about the effect, the other because his answer could not be clearly interpreted.

claim that there is no conscious perception without attention. The consistency of the theoretical conclusions that can be drawn from these two radically different paradigms is reassuring. Whether the unexpected object is flashed for 200 ms in an otherwise empty display or it moves dynamically across a natural scene for 5 s, observers are unlikely to notice it if attention is otherwise engaged. This consistency suggests that the results of computer-based studies of inattention blindness can and do generalize to situations closer to our real-world experiences.

The results of our experiments also call to mind recent findings from research on change blindness. Many studies of change blindness focus on simple displays of letters or dots to determine how little information is preserved from one view to the next. More recently, change-blindness research has moved from using simple displays of letters, dots, and words (eg Pashler 1988; Phillips 1974) to more complex, naturalistic displays for which more information is available for selection (see Simons and Levin 1997 for a review). Given the simplicity and relative meaninglessness of the simple displays, the generalizability of the results to more naturalistic viewing conditions was not certain (see Simons 2000, for discussion). The recent thrust of work on change blindness has been to examine whether the inferences drawn from work with simple displays will hold for more natural displays. One dramatic demonstration was at least partly responsible for this move toward increased naturalism. When viewing photographs of natural scenes in preparation for a memory test, people missed large, meaningful changes that occurred during eye movements (Grimes 1996). For example, observers often failed to notice when two people in a photograph exchanged hats or even when they exchanged heads. These findings have been replicated in subsequent work on saccade-contingent changes (Currie et al 1995; Henderson and Hollingworth, *in press*; McConkie and Currie 1996).

This change blindness for natural scenes has been extended to a number of other paradigms. For example, when an original and modified image are presented in rapid alternation with a blank screen interposed between them, observers have great difficulty detecting changes (Rensink et al 1997). This 'flicker' technique shows that change blindness is not limited to saccade-contingent changes. In the case of saccade-contingent changes, the blur on the retina caused by the eye movement itself leads to suppression of visual processing during the change, thereby preventing detection of any local transients. The blank screen in the flicker study has essentially the same effect, producing a global change signal that prevents detection of the local one caused by the change. Similar change blindness has been shown for changes across cuts or pans in motion pictures (Levin and Simons 1997; Simons 1996), eye blinks (O'Regan et al 2000), and 'mud splashes' (O'Regan et al 1999). As noted earlier, even when one conversation partner is replaced by a different person during a brief interruption, observers often fail to notice the change (Simons and Levin 1998).

As in studies of inattention blindness, the likelihood of change detection depends on the focus of attention. In studies of inattention blindness, when observers are attending to another object or event, they are less likely to notice the unexpected event. In studies of change detection, people are better able to report changes to attended than unattended objects. For example, people are faster to detect changes in the flicker paradigm when the changed object is of central interest in the scene (Rensink et al 1997). Central objects are more likely to garner attentional resources, and if we have a limited capacity for holding information across views, changes to objects that receive more effortful processing are more likely to be detected (see Rensink 2000; Scholl 2000). Just as we often fail to perceive unexpected events, we often fail to notice unexpected changes to the visual details of our environment—in both cases, this applies even when attention is focused on the area of the event or change.

Although the theoretical conclusions drawn from real-world studies are not altogether different from those derived from work with simple displays, they do show that change

blindness is a general property of the visual system and that it applies to almost all aspects of visual processing. We apparently do not retain a detailed visual representation of our surroundings from one view to the next, even for displays with all the richness of natural scenes. Similarly, studies of sustained inattention blindness suggest that we fail to perceive unexpected objects even under naturalistic viewing conditions.

The results of our studies of sustained inattention blindness, however, do contrast in an interesting way with those of one recent change-blindness study (Simons et al, in preparation). In that experiment, a female experimenter dressed in athletic clothing and carrying a basketball approaches a passerby in public and asks directions to a gym. During this interaction, a crowd of confederates walked between the two and surreptitiously took the basketball away. When asked if they noticed anything changed or anything different about her appearance, a minority of observers reported noticing that the basketball was gone. But when asked a follow-up question specifically referring to the basketball, most of the remaining observers 'remembered' the basketball and were able to describe its unusual coloring. Thus, a visual change can be encoded but not explicitly reported until a specific retrieval cue is provided. In the experiment reported here, however, not one of the eighty-eight nonnoticers 'remembered' the Gorilla or Umbrella-Woman events when specifically asked about it, and several did not believe that the event had happened until the videotape was replayed for them.

While there are several important differences between these paradigms that could account for this difference in behavior, they share the feature that a condition of inattention was created (by the conversation in the basketball disappearance study or by the monitoring task here) that apparently prevented many observers from becoming aware of a salient visual change. Perhaps the crucial difference is that whereas the conversation simply reduced the observer's attention by drawing it away from the critical object, the monitoring task in this study required observers to attend to one event *while ignoring another* that was happening in the same region of space. This 'directed ignoring' could inhibit perception of not just the ignored event but of *all* unattended events, thereby preventing the formation of an explicit memory trace. Whether inattention blindness occurs because the target is similar to the intentionally ignored items or different from the attended items is an open question that would be relatively difficult to explore by using video-based displays but could be explored by using more controlled computer-based tasks.

One alternative interpretation of our findings is that subjects did consciously perceive the unexpected object, however briefly, but immediately forgot they had seen it (Wolfe 1999). Although this inattention-amnesia explanation can in principle account for our findings, it seems less plausible that the inattention-blindness account for a number of reasons. First, detecting unusual objects or events would be a useful function for a visual system to have; immediately forgetting them would defeat this purpose. This is especially true for a prolonged, dynamic event. Given that the unexpected object in our experiments was available for further examination (something that was less true of earlier studies with briefly flashed objects), we might expect observers to try to verify their percept in these studies, thereby leading to a preserved representation. Furthermore, if observers did consciously perceive and then forget the gorilla, they presumably would not be particularly surprised when asked if there had been a gorilla in the display. Yet, observers in our study were consistently surprised when they viewed the display a second time, some even exclaiming, "I missed *that*!?" It seems more parsimonious to assume that observers were never aware of the unexpected object than to assume that they saw a gorilla, then forgot about it, and then were shocked to see it when told to look for it. Last, as noted earlier, Becklen and Cervone (1983) found no improvement in noticing when the video was stopped immediately after the unexpected event rather than several seconds later. However, finding a direct test to distinguish between never perceiving an object and

immediately forgetting it will be difficult because the inattentional-amnesia proponents could always argue that the memory test came too late. Thus, it may not be possible to distinguish empirically between the amnesia and the blindness explanations.

Although our findings suggest that unexpected events are often overlooked, the question of whether they leave an implicit trace remains open. Unnoticed stimuli in the static-inattentional-blindness paradigm can lead to priming effects (Mack and Rock 1998). However, those experiments did not require subjects to ignore anything. Neisser and colleagues found that subjects under the conditions we have described as directed ignoring were no more likely to select the unexpected object in a two-alternative forced-choice recognition test than were other subjects when asked to report it directly (Neisser and Rooney 1982, as cited in Becklen and Cervone 1983). However, forced choice may not be as sensitive as other implicit memory tests. Future research should explore the issues of preserved representations and directed ignoring within the sustained-inattentional-blindness paradigm we have reintroduced here.

Acknowledgements. Many thanks to all of the people who helped with filming the videos or collecting the data for this study: Jennifer Shephard (Umbrella Woman), Elisa Cheng (Gorilla), Judith Danovitch, Steve Most, Alex Wong (White team), Amy DeIpolyi, Jason Jay, Megan White (Black team), Dan Ellard, Samantha Glass, Jeremy Gray, Sara Greene, Anya Hernandez, Orville Jackson, Latanya James, David Marx, Steve Mitroff, Carolyn Racine, Kathy Richards, Chris Russell, Laurie Santos, Steve Stose, Ojas Tejani, Dan Tristan, Amy Wiseman, Leah Wittenberg, and Amir Zarrinpar (data collection, in alphabetical order). Thanks also to M J Wraga for suggesting the first part of the title of this article, to Brian Scholl for discovering the Bálint quotation, to Dick Neisser for inspiration and for giving permission to produce figure 1, to Larry Taylor for helping to avoid collisions during the filming of the videos, and especially to Jerry Kagan for lending us his gorilla suit. Additional thanks to Steve Most, Brian Scholl, two anonymous reviewers, and everyone else who commented on earlier versions of this manuscript and presentations of these results.

Miniaturized and abbreviated versions of the videos used in this study are available, in QuickTime format, via the Internet at: <http://coglab.wjh.harvard.edu/gorilla/index.html>. No animals were harmed during the making of the videos.

References

- Bahrick L E, Walker A S, Neisser U, 1981 "Selective looking by infants" *Cognitive Psychology* **13** 377–390
- Becklen R, Cervone D, 1983 "Selective looking and the noticing of unexpected events" *Memory and Cognition* **11** 601–608
- Cherry E C, 1953 "Some experiments upon the recognition of speech, with one and with two ears" *Journal of the Acoustical Society of America* **25** 975–979
- Chun M M, Jiang Y, 1998 "Contextual cueing: Implicit learning and memory of visual context guides spatial attention" *Cognitive Psychology* **36** 28–71
- Currie C, McConkie G W, Carlson-Radvansky L A, Irwin D E, 1995 "Maintaining visual stability across saccades: Role of the saccade target object" technical report UIUC-BI-HPP-95-01, Beckman Institute, University of Illinois
- Grimes J, 1996 "On the failure to detect changes in scenes across saccades", in *Perception (Vancouver Studies in Cognitive Science)* Ed. K Akins, volume 2 (New York: Oxford University Press) pp 89–110
- Haines R F, 1989 "A breakdown in simultaneous information processing", in *Presbyopia Research: From Molecular Biology to Visual Adaptation* Eds G Obrecht, L W Stark (New York: Plenum) pp 171–175
- Henderson J M, Hollingworth A, in press "The role of fixation position in detecting scene changes across saccades" *Psychological Science*
- Husain M, Stein J, 1988 "Rezső Bálint and his most celebrated case" *Archives of Neurology* **45** 89–93
- Levin D T, Momen N, Drivdahl S B, Simons D J, 2000 "Change blindness blindness: The meta-cognitive error of overestimating change-detection ability" *Visual Cognition* (in press)
- Levin D T, Simons D J, 1997 "Failure to detect changes to attended objects in motion pictures" *Psychonomic Bulletin and Review* **4** 501–506
- Littman D, Becklen R, 1976 "Selective looking with minimal eye movements" *Perception & Psychophysics* **20** 77–79
- Mack A, Rock I, 1998 *Inattentional Blindness* (Cambridge, MA: MIT Press)

- Mack A, Tang B, Tuma R, Kahn S, 1992 "Perceptual organization and attention" *Cognitive Psychology* **24** 475–501
- 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
- McConkie G W, Zola D, 1979 "Is visual information integrated across successive fixations in reading?" *Perception & Psychophysics* **25** 221–224
- Moore C M, Egeth H, 1997 "Perception without attention: Evidence of grouping under conditions of inattention" *Journal of Experimental Psychology: Human Perception and Performance* **23** 339–352
- Moray N, 1959 "Attention in dichotic listening: Affective cues and the influence of instructions" *Quarterly Journal of Experimental Psychology* **11** 56–60
- Neisser U, 1979 "The control of information pickup in selective looking", in *Perception and its Development: A Tribute to Eleanor J Gibson* Ed. A D Pick (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 201–219
- Neisser U, Becklen R, 1975 "Selective looking: Attending to visually specified events" *Cognitive Psychology* **7** 480–494
- Neumann O, Heijden A H C van der, Allport D A, 1986 "Visual selective attention: Introductory remarks" *Psychological Research* **48** 185–188
- Newby E A, Rock I, 1998 "Inattention blindness as a function of proximity to the focus of attention" *Perception* **27** 1025–1040
- O'Regan J K, Deubel H, Clark J J, Rensink R A, 2000 "Picture changes during blinks: Looking without seeing and seeing without looking" *Visual Cognition* (in press)
- O'Regan J K, Rensink R A, Clark J J, 1999 "Change-blindness as a result of 'mudsplashes'" *Nature (London)* **398** 34
- Pashler H, 1988 "Familiarity and visual change detection" *Perception & Psychophysics* **44** 369–378
- Phillips W A, 1974 "On the distinction between sensory storage and short-term visual memory" *Perception & Psychophysics* **16** 283–290
- Rensink R A, 2000 "Visual search for change: A probe into the nature of attentional processing" *Visual Cognition* (in press)
- Rensink R A, O'Regan J K, Clark J J, 1997 "To see or not to see: The need for attention to perceive changes in scenes" *Psychological Science* **8** 368–373
- Rock I, Linnett C M, Grant P, Mack A, 1992 "Perception without attention: Results of a new method" *Cognitive Psychology* **24** 502–534
- Rooney P, Boyce C, Neisser U, 1981 "A developmental study of noticing unexpected events" unpublished manuscript
- Rubin N, Hua X L, 1998 "Perceiving occluded objects under conditions of inattention" *Investigative Ophthalmology & Visual Science* **39**(4) S1113
- Scholl B J, 2000 "Attenuated change blindness for exogenously attended items in a flicker paradigm" *Visual Cognition* (in press)
- Silverman M, Mack A, 1997 "Priming by iconic images" *Investigative Ophthalmology & Visual Science* **38**(4) S963
- Simons D J, 1996 "In sight, out of mind: When object representations fail" *Psychological Science* **7** 301–305
- Simons D J, 2000 "Current approaches to change blindness" *Visual Cognition* (in press)
- Simons D J, Chabris C F, Levin D T, (in preparation) "Change blindness is not caused by later events overwriting earlier ones in visual short term memory"
- Simons D J, Levin D T, 1997 "Change blindness" *Trends in Cognitive Sciences* **1** 261–267
- Simons D J, Levin D T, 1998 "Failure to detect changes to people in a real-world interaction" *Psychonomic Bulletin and Review* **5** 644–649
- Stoffregen T A, Baldwin C A, Flynn S B, 1993 "Noticing of unexpected events by adults with and without mental retardation" *American Journal on Mental Retardation* **98** 273–284
- Stoffregen T A, Becklen R C, 1989 "Dual attention to dynamically structured naturalistic events" *Perceptual and Motor Skills* **69** 1187–1201
- Treisman A, 1964 "Monitoring and storage of irrelevant messages in selective attention" *Journal of Verbal Learning and Verbal Behavior* **3** 449–459
- Williams P, Simons D J, 2000 "Detecting changes in novel, complex three-dimensional objects" *Visual Cognition* (in press)
- Wolfe J M, 1999 "Inattentional amnesia", in *Fleeting Memories: Cognition of Brief Visual Stimuli* Ed. V Coltheart (Cambridge, MA: MIT Press) pp 71–94