Inattentional blindness and augmented-vision displays: Effects of cartoon-like filtering and attended scene

Henry L. Apfelbaum, **Doris H. Apfelbaum**, **Russell L. Woods**, and **Eli Peli**
Schepens Eye Research Institute and Harvard Medical School, Boston, MA, USA

Abstract

Augmented-vision devices that we are developing to aid people with low vision (impaired vision) employ *vision multiplexing* – the simultaneous presentation of two different views to one or both eyes. This approach enables compensation for vision deficits without depriving the wearers of their normal views of the scene. Ideally, wearers would make use of the simultaneous views to alert them to potential mobility hazards, without a need to divide attention consciously. *Inattentional blindness*, the frequent inability to notice otherwise-obvious events in one scene while paying attention to another, overlapping, scene, works against that sort of augmentation, so we are investigating ways to mitigate it. In this study we filtered the augmented view, creating cartoon-like representations, to make it easier to detect significant features in that view and to minimise interference with the normal view. We reproduced a classic inattentional blindness experiment to evaluate the effect, and found that, surprisingly, edge filtering had no detectable effect – positive or negative – on the noticing of unexpected events in the unattended scene. We then modified the experiment to determine if the inattentional blindness was due to the confusion of overlaid views or simply a matter of attention, and found the latter to be the case.

Keywords

Inattentional blindness; augmented vision; low-vision aids; cartooning; edge filtering; unexpected events

Introduction

Seeing two different scenes superimposed on one another, with both in view simultaneously in both eyes, rarely occurs in the natural world, but it is increasingly common in devices which augment vision. The human visual system has had little need or opportunity to adapt to displays of that sort. Seeing one's reflection on a near-transparent surface, as can happen when viewing fish in a stream, may be the only sort of visual overlap encountered during human evolution. By comparison, head-up displays and augmented-vision devices in use by the military and the civilian sector are proliferating (Martin-Emerson and Wickens, 1997; Stevens *et al.*, 1998), but not without problems (e.g., Haines, 1991).

Many devices we have developed to aid people with impaired vision make use of vision multiplexing; the simultaneous presentation of two different views to one or both eyes (Peli, 2001; Luo and Peli, 2006; Peli, 2007; Peli *et al.*, 2007). In particular, augmented-vision approaches for patients with peripheral visual field loss (as due to retinitis pigmentosa or glaucoma), present the augmented view in the primary position of gaze. For example, spectacles fitted with a small video camera and a display that is viewed via a see-through beam

Correspondence and reprint requests to: Henry Apfelbaum, Schepens Eye Research Institute, 20 Staniford Street, Boston, MA 02114, USA, Tel: 011 617 912 2588; Fax: 011 617 912 0112, E-mail address: Henry.Apfelbaum@Schepens.Harvard.edu.

splitter in one lens (Vargas-Martin and Peli, 2002; Bowers *et al.*, 2004) are used to provide a minified wide-angle view within the limited visual field of patients with severe peripheral visual field loss. One such device provides an 80° field of view within a display that subtends 16°. The minified display is edge-filtered to create cartoon-like image contours, in order to reduce interference with the see-through view and to make it easier to isolate features important for mobility (Figure 1). In the figure, the person and trash container can be seen in the minified view but would be missed if not scanning the normal view. We need to know if the visual system can handle this information effectively.

In a classic psychophysical experiment on "selective looking", Neisser and Becklen ("N&B", 1975), identified the phenomenon of *inattentional blindness* (IB) — the apparent inability to notice significant but unexpected events (UEs) in an unattended scene when attention is fixed on another scene, even though they both are in the same area of the visual field. They showed normally-sighted subjects videos of two different games at the same time; a ball game with three men running in a circle and tossing a basketball to one another, and a "hand game", with the outstretched hands and forearms of two players who alternated attempting to slap the other's hands. They kept their subjects' attention focused on one of the games by assigning a task that demanded attention (pressing a switch at each ball-toss or hand-slap attempt), and occasionally introduced UEs in the other game, such as replacing the male ball-game players, one by one, with females, or having the hand-slappers stop briefly to shake hands. A surprisingly large number of subjects failed to notice some of the UEs.

The phenomenon has proved to be extremely robust, and was given the name IB by Mack and Rock, who devoted an entire book to the subject (Mack and Rock, 1998). They used synthetic displays to probe features which led to IB, and showed, for example, that UEs could be missed even when presented foveally. Others have also studied IB, sometimes with synthetic displays and sometimes with natural scenes in a manner similar to N&B, often to probe just how similar or co-located the action of the overlapping scenes could be and still induce IB. They have shown that action can be co-located and very similar (Becklen and Cervone, 1983; Most *et al.*, 2000; Most *et al.*, 2001; Koivisto *et al.*, 2004). Simons and Chabris (1999) showed that IB persisted even if two instances of the same game physically coexisted in one scene, rather than in just the semi-transparent views of overlaid videos.

IB is an undesirable phenomenon if it occurs when patients are using our vision multiplexing aids, yet all studies would seem to predict its likelihood. For example, we want pedestrians with severe peripheral visual field loss, using our augmented-vision spectacles while waiting to cross a road, to notice an approaching car, even though the car is only visible in the wide-angle view provided by the minified display and they are paying attention to the see-through view. However, IB reduces the likelihood of such events being noticed. Hence we are seeking ways to mitigate IB.

The efficacy of our devices is also potentially constrained by a complementary phenomenon, *change blindness* (Simons and Levin, 1997; Rensink, 2000). Viewers often fail to notice significant changes in a scene when attention is distracted or interrupted. Ideally, if a wearer can benefit from overlapping views without needing to alternate attention between them, the devices would be more effective, and this would argue for having both scenes in view simultaneously, and this, too, adds to the desire to mitigate IB.

In Experiment 1, we tested the hypothesis that cartooning would affect IB. We reproduced the relevant aspects of the Neisser and Becklen ("N&B") study, and then treated one or both of the overlapping views with an electronic edge filter that reduced the full-colour video to cartoon-like outlines of the features in the scene. We expected that cartooning would cause less interference between the views and might mitigate IB, while drawing attention to important

features in the cartooned scene. This experiment was not intended to be directly predictive of the effects of cartooning in our augmented vision devices, which deal with likely events and show two views of the same scene. Rather, it established a baseline test of the effects of cartooning that could be compared directly with the N&B results.

To our knowledge, no one had tested UEs in a naturalistic (not symbolic) *attended* scene while otherwise maintaining the overlapping views. Thus we raised the question of whether it was the overlapping that caused IB or inattention itself. We realised that it would be simple to test this with videos we had already prepared and used in Experiment 1, by simply changing the instructions so that the subject would attend to the action of the scene that contained the UE, and that is what we did in Experiment 2.

Experiment 1: Effect of cartoon-like edge filtering

Cartooning one of the overlapping scenes in an N&B-like study might affect the ease with which the scenes can be distinguished from one another, while preserving and perhaps emphasizing salient features of the cartooned view. We hypothesized that this was likely to affect the degree of IB engendered by the overlapping presentations. Experiment 1 tested that hypothesis.

Methods

Videos

A Canon ZR10 miniDV camcorder was used to tape the source videos used in this experiment. Excerpts from each of the videos described in the figures can be played by clicking on the images.

Games

As in N&B (Neisser and Becklen, 1975), scenes of three young men tossing a basketball and running in a circle were videotaped (Figure 2a). Occasionally they feinted passes, and frequently dribbled without passing. Passes were either direct or via a single bounce. This basic scene was recorded 6 times (6 "takes"), and the 4 takes with smoothest action were selected for use as attended scenes in the experiment (and as unattended scenes in the trials that did not include UEs). In all takes, the ball was passed 30 times during 60 s of play.

Similarly, the hand game was taped 6 times and the best 4 takes were selected for use as attended scenes. (Figure 2b). One player (the slapper) had his palms up, and the other player had his palm resting on the slapper's palms. The slapper quickly turned over one or both of his hands and attempted to slap the other player. If he missed, the slapper and target players exchanged roles. In all takes there were 30 slap attempts during 60 s of play. Occasionally the slapper would feint a slap, not turning his hands over completely. Feints did not count in the thirty slap attempts.

Unexpected Events

Five different ball game scenes and three different hand game scenes with unexpected events were taped for use as unattended scenes in the experiment. Each UE scene contained approximately 30 passes or hand slap attempts. Figure 3 shows each of the hand game UE scenes, and Figure 4 shows the ball game scenes.

Cartooning

A video edge-filtering device was developed to our specifications by DigiVision, Inc. (San Diego, CA), as a variant of their ValueVision filter. The filter processes the S video luminance

channel. Each take was transferred from the computer in DV format to a camcorder, played from that camcorder as S Video through the filter to a second camcorder, and then recaptured, thus avoiding deinterlacing that the computer's S Video input always performs. Deinterlacing artefacts are particularly objectionable with edge-filtered video motion.

Test recordings were made to establish filter contrast and threshold settings that seemed to provide the clearest edges with a minimum of noise. The filter was used in its bipolar binary mode. The bipolar mode is unique to this version of the filter (Peli, 2002; 2004). The nominal off-edge output of the filter is grey (~53 IRE, ~0.4v). Detected edges are represented by a both a positive-going and negative-going transition from the nominal value. In analogue mode, the magnitude of the swing is proportional to the strength of the edge. In bipolar mode, the swing is to full white (100 IRE, 0.714v) and full black (7.5 IRE, ~0.05v). Bipolar mode is especially effective, as it ensures that edges will show against both light and dark backgrounds (Figure 5).

Presentations

A *presentation* is a 75-s video segment shown to a subject during a trial. Each presentation started with a 15-second lead-in segment with synchronization signals at 5 and 10 s, followed by 60 s of game play. The sync signals were single bounces, if the ball game was to be attended, or single finger snaps, if the hand game was to be attended. The play period had just the attended game if only one game was to be shown (during the practice trials described below), or a superimposition of a hand game scene and a ball game scene. The attended and unattended scenes could each have been in full colour or cartooned, giving 4 possible cartooning treatment conditions for a presentation (Figure 5, a, c, e, and f).

The full-colour (unfiltered) videos and the raw edge-filtered videos were processed using Adobe Premiere 6.5 to combine them into presentations. (We have since developed a video mixer that can perform the superimposition in real time.)

If both scenes of a presentation were to be in full colour (neither cartoon-like), they were simply combined with the Premiere overlaying track set at 50% transparency. If both were to be cartooned, the black/grey/white edge-filtered video was converted to full black (off edge) and white (edge) via appropriate track threshold settings, and the views were merged.

If just one of the scenes was to be shown in cartoon-like form, intermediate tracks were used to create masks that let the lower track through wherever the upper track was grey, and otherwise show the white or black of the upper track. Contrast was set high in the masks to preserve crisp edges.

In all, the experiment design called for the creation of 145 presentations, as described below. To avoid introducing compression artefacts in the cartooned views, uncompressed DV format was preserved throughout.

Subjects

The experiment was balanced along several dimensions, as described below, requiring 36 subjects, one per session. Subjects were recruited from local campuses and by on-line ads on craigslist.com and Biotrax.com. Thirty-eight subjects were recruited, but two were subsequently excluded (one due to age criteria and the other due to a post-experiment admission of familiarity with IB). The remaining 36 (11 males) met the inclusion criteria, being between 18 and 40 years old (20–34 years), with normal, or corrected to normal, eyesight (20/30 or better). Subjects signed Schepens Institutional Review Board-approved consent forms and were paid $10 for their session. Each session took an hour or less.

Balancing

As described in Appendix A, a moderately complex balancing scheme was employed. Each subject viewed all 4 possible cartooned/not cartooned treatment combinations for each game (8 scored presentations in total). Since no subject should see a UE scene more than once while testing for IB, it was not possible to show a subject all UEs with all possible cartooning treatments per UE. Rather, each UE within a game was paired with a different treatment combination. Each of the UEs (3 per game) was shown in 6 of the 8 presentations, while each subject viewed one presentation per game with no UE (2 in total). The presentations with both scenes edge-filtered were those with no UE, as this format was not of interest for the augmented-vision devices under development. The presentations without a UE ensured that the experimenter did not know which presentations contained UEs. To avoid familiarisation with the attended game events, each attended showing of a game used a different take (and hence the 4 takes of each game).

A power analysis of this approach found that the experiment would have a 24% chance of finding a small difference among the detection rates of UEs paired with each cartooning treatment, and a 97% chance of detecting a medium difference (where "small" and "medium" are attributed to effect size $w = 0.10$ and 0.20, respectively, in Cohen, 1988). This was considered adequate, since small differences would not be of practical relevance for the vision aids under consideration.

The order in which UEs were presented and their treatment might affect the number detected, due to possible priming effects. Therefore, the order of UEs and the pairing of UEs with cartooning treatments were balanced across subjects. It also might be easier to detect some UEs when paired with particular attended game takes. As it was not feasible to test all combinations of pairings and orders, we used a combination of balancing, partial balancing, and randomization to reduce the number of subjects needed.

Physical setup

The subject and experimenter sat facing each other across a table. The experimenter had a computer monitor, keyboard and mouse, while the subject had a 15″ diagonal TV monitor and a mouse. Neither could see the other's screen. This setup, together with the pseudo-randomized order of presentations, ensured that the experimenter would not know which trials included UEs and could not give subconscious cues to the subject. The TV monitor was adjusted to be at the viewing distance the subject preferred, generally about 1 m. The experimenter's workstation provided the prompts read to the subject and collected the responses. The subject used a mouse click to indicate that a pass or slap attempt occurred in the attended scene.

Session Procedures

Prolonging naïveté while detecting detections

There is evidence that IB experiments can be conducted successfully even when the subject expects "unexpected" events, if their onset occurs when the subject is known to be fixating on an event in a distractor task and their duration is short (Rensink, 2005). We did not rely on such timing. In this study, the UEs were long enough to easily be noticed when scanning, so it was important to ensure that they were truly unexpected. To mask the true nature of the experiment, we told subjects that our purpose was to determine how well the visual system can deal with overlapped images, and to see if cartooning an image makes a difference. (We used our augmented-vision spectacles as an example.) We explained that we would be showing them overlapped videos and asking them to pay attention to the action in one of them, and afterwards we would be asking them questions about the difficulty of the task. No mention was made of UEs.

Having thus set the scene, we then casually asked if they had ever seen videos overlapped in that way, on TV, for instance, or in psychology classes. If they replied that they had not, we considered them to be suitably naïve to IB experiments. Had they seen, for instance, the widely-shown video clip from the Simons and Chabris (1999) study, it was likely that these questions would have brought it to mind. No subject was excluded that way, although one subject later remembered seeing the Simons and Chabris video.

The questions asked after each presentation was shown were also designed to avoid alerting the subjects to the existence of the UEs. The subjects were read multiple-choice questions asking them to rate how difficult the task was, and to identify any particularly hard parts. The experimenter acted very interested in that feedback, and carefully typed in all comments the subjects made.

The question intended to determine if a UE was detected was: "Was there anything worth noting in the background video that was distracting or interfered with following the game?"

To further prolong naïveté, subjects were told that the experiment was still undergoing pilot tests, and depended on computer software to randomly select the videos they would be shown. When a subject first identified a UE, the experimenter expressed surprise, asking just what had been seen. Then she "realized" that a video from another study must have gotten mixed in, and said that shouldn't affect the outcome of this study, so they could continue. When a second UE was detected, the experimenter's surprise turned to irritation with the programmer, but again said that continuing would be okay. By the third detection, the experimenter just acted resigned to the appearance of the "wrong" videos, and pressed on.

Trials

Each session used a different combination and ordering of presentations and included up to 26 trials (Appendix B, Tables B1–3). The first four trials introduced the subject to the video treatments and task. The next eight (trials 5–12) were used in the analyses. Six of those 8 trials included UEs.

In each of those trials, the subject was given instructions about the task, shown a presentation, and then asked the follow-up questions described above. The subject was instructed to click the mouse once for each of the synchronization bounces or snaps, and then once for each event (toss or slap attempt) in the subsequent game. The subject was told that the first several trials were just for practice. Every four trials the experimenter asked if the subject was tiring and could use a break. No subject found a break necessary.

The experimenter entered the responses to each of the multiple-choice questions that followed each presentation, together with any comments the subjects made. For the interference question, the experimenter did not use the subjects' answers directly. Rather, if a subject responded that there was interference or a distraction, the experimenter asked what that was, and tried to judge if the response indicated that the subject had noticed one of the UEs, and would enter that information. So, for example, if subjects said that the motion of the hand-slappers' arms made it difficult to follow the basketball when it was tossed behind the hands, that was not considered to be an event detection, but if subjects said that they were distracted when the hand-slappers stopped to play rock-paper-scissors, the experimenter scored it as a detection of the choose-up event.

After trial 12, increasingly revelatory questions and replays were employed to determine if the subject had failed to acknowledge some events that had indeed been detected, and the detection scores were adjusted accordingly. Appendix B gives the details of the trial order and the follow-up questions and replays.

Session control software

A Microsoft Access/VBA application was used to control trial order and content, provide the prompts and questions that the experimenter read to the subject, play the presentations, record the response times, and record the multiple-choice responses and free-form comments, and store the data for later analyses. It also checked that the user responded properly to the synchronization events in each trial's lead-in, and restarted the trial if the events were missed or a false alarm occurred.

Measures

For each subject, each of the 6 UEs was scored as detected or undetected, as described above. Two additional measures were derived for each of the experimental trials 5 through 12: average response time to attended game events, and accuracy of responding to those events (hit rate). Response time differences could indicate differences in degree of difficulty of the attended task under the different conditions, and a low hit rate would indicate a lack of attention to the distractor task.

The time of occurrence of each mouse click in response to an attended game event was recorded. The click was scored as a hit if it occurred within a window 0.5 s before to 0.5 s after a game event plus the subject's average response time to hits in that trial. The average response time used in the statistical analyses is the average difference between the actual time of the hit clicks in a trial and the corresponding game events of that trial (the middle of the events, not the onset).

Statistical techniques

Chi-squared (χ^2) analysis and Cochran's Q (Q) were used to judge statistical significance of event detections, while repeated measures ANOVA (F) and the Student t test (t) were used for response time analyses, and Friedman analysis of variance by ranks (Fr) was used for hit accuracy statistics. Significance levels *(p)* associated with the Cochran's Q tests assume that χ^2 approximates Q, as the sample sizes were sufficient to justify that assumption. The Wilcoxon signed rank test (Z) was used to judge significance of accuracy differences by game. The significance of differences between proportions (or probabilities) was assessed using Method 10 from Newcombe (1998).

A $p \leq 0.05$ was considered to have statistical significance.

Results

Unexpected event detection

UEs were detected in 123 (57%) of the trials in which they were shown. Figure 6 shows the number of subjects who detected a given number of UEs. Table 1 shows detections by UE scene and cartooning treatment combination.

Cartooning did not have a significant effect on detection rate ($\chi^2_{(2,16)} = 0.79$, $p = 0.67$). Analysed separately by game, cartooning had no significant effect on detection of UEs in the ball game scenes (Q = 2.15, $p = 0.34$) or hand game scenes (Q = 0.09, $p = 0.96$). Nor was cartooning significant when analysed by number of times (0, 1, or 2) that a UE with a particular cartooning treatment was detected within each session ($\chi^2_{(4,108)} = 1.93$, $p = 0.75$).

As can be seen in Table 1, some UEs were detected significantly more frequently than others ($\chi^2_{(5,216)} = 54.8$, $p < 0.001$).

Effects on response time

Edge treatment had a modest, but statistically significant effect on response times ($F_{(3, 105)} = 5.52$, $p < 0.001$). Response time to the attended task decreased if the unattended task was cartooned, and increased if the attended task was cartooned (Table 2).

Response time also differed by game, ($t_{(35)} = 6.2$, $p < 0.001$). Average response time to ball game tosses was 462 ms, while average response time to hand-slap attempts was 563 ms.

Not all takes of a game were equally easy to follow, as evidenced by significant response time variation by take ($F_{(3,105)} = 82.2$ and 22.7, respectively, for the ball game and hand game, $p < 0.001$, Figure 7).

Even so, attended game response times did not vary significantly when analysed by the paired UE scene ($F_{(2,70)} = 0.35$, $p = 0.71$, for hand game response time per ball game UE scene and $F_{(2,70)} = 0.43$, $p = 0.65$, for ball game response time per hand game UE scene).

Effects on accuracy

Cartooning had no significant effect on hit rate ($Fr = 1.18$, $p = 0.76$).

Although the numbers of hits for the attended game events per trial were all near ceiling (averaging 97%), 35 of the 36 sessions had an accuracy difference between games, and the difference was statistically significant ($Z_{34} = 4.5$, $p < 0.001$), with lower accuracy when attending the ball game (95%) than when attending the hand game (98%). Some subjects mentioned that the ball game was harder to follow because occasionally a player would come between the ball and the camera, so that a toss could only be surmised after the fact. A response to a hidden ball could be delayed long enough to be scored as a miss rather than a hit, and thus reduce the average accuracy and not increase the average hit response time. Delayed responses of that sort would also be scored as false alarms as well as misses. The false alarm rate of the ball game also included true false alarms, in which the subject clicked when there was no toss (as occasionally happened, for instance, when a dribble was taken as a toss). The false alarm rate of the ball game (5%) was slightly higher than that of the hand game (4%), presumably due to the effectiveness of feints in the hand game, and the difference was statistically significant ($Z_{29} = 3.0$, $p = 0.003$).

Accuracy varied significantly across the 4 attended ball game takes ($Fr = 10.1$, $p = 0.02$), although the actual variation was not large, with the hit rate of all takes falling between 94% and 97%. No such effect was found for the hand game ($Fr = 2.3$, $p = 0.52$).

Attentiveness

Subjects who noticed many of the UEs did not sacrifice attended task performance to achieve it. There was no significant correlation between number of UEs noticed and response time or accuracy (Pearson $R^2 < 0.06$, $p > 0.6$).

Effect of attended take

The umbrella woman's detection rate of 53% was closest to the overall average rate of 57%, but the rate varied over a range of 14 to 83%, depending on the hand game take she was paired with. No other UE had such a large range of detections, and the effect of attended take on the umbrella woman's detection rate was significant ($\chi^2_{(3, 36)} = 11.7$, $p = 0.008$). The corresponding significance levels for the lost ball, ball toss, choose-up, juggler and handshake events were 0.09, 0.48, 0.89, 0.95 and 0.95, respectively.

Order effects (priming)

Since each subject was shown more than one UE (i.e., 6), the possibility exists that subjects were more likely to notice UEs once they had already detected one, and that could render detection rate statistics after the first noticed event relatively meaningless. In fact, we did find a small priming effect. The probability of detecting a UE in a trial that occurred before any UE had been detected in a session was 0.46, while the probability of detecting a UE after at least one had been detected was 0.62. The difference between these two probabilities is significant ($p = 0.02$). Since the order of presentation of UEs was balanced across sessions, the difference is likely due to priming, not any skew in the distribution of event difficulty.

Dividing each session's UE trials into 3 phases defined as 1) trials before the first detection, 2) the first detection trial and 3) trials after the first detection, no significant priming effect of detection on attended task accuracy or response times was found ($F_r < 0.01$, $p > 0.99$, for hit accuracy by detection phase, and $F_{(2, 36)} = 0.04$, $p > 0.95$, for response time by detection phase).

Discussion

We applied cartoon-like edge filtering to one or both of the overlaid scenes in an experiment modelled on the classic N&B study that introduced IB, and our results are in essential agreement with theirs.

Cartooning would seem to be a fairly radical treatment. Finding no effect of cartooning on UE detection was therefore surprising. It neither increased nor decreased the event detection rate, regardless of whether the attended or unattended scene was cartooned. Our failure to detect an effect might be due to masking by the strong effect of UE scene, but as noted above, the experiment had power enough to detect any effect of a magnitude interesting enough for our application to low-vision aids. While perceptual and cognitive load are known factors in inducing IB (e.g., Lavie *et al.*, 2004; Todd *et al.*, 2005; Lavie, 2006; Cartwright-Finch and Lavie, 2007), they are essentially constant across the attended scenes of each game. We can only conclude, as others have in studies with different manipulations, that the contextual relationship of the events to the attended scene plays an important role in inducing or mitigating IB, and at least in this case, is a stronger factor than visual parameters or location. For example, see the unpublished report by Becklen and Neisser described in (Neisser, 1979) and also (Becklen and Cervone, 1983; Simons and Chabris, 1999).

We included hit accuracy and response time measures primarily for control purposes, to check that subjects with high UE detection scores were not achieving them at the expense of the attended distractor task, and we feel that was borne out. The response time data, did, however, exhibit a statistically significant effect of cartooning on the attended tasks, as non-filtered (full-colour) presentations seemed to be a little easier to follow. The absolute difference, however, was small, and not likely to be of consequence in our applications.

It is interesting to speculate on what aspects of the UEs we used led to higher or lower relative detection rates, as this might be of use in the design of future experiments, and in the design of low-vision aids. The lost ball event in the ball game was detected least, with its detection rate of 19% compared to the average detection rate of 57%. The players were very good at mimicking ordinary play, even though they were playing without the ball, and even the onset and termination of the event with the loss and recapture of the ball occurred very smoothly, near the edge of the frame. So it was simply the absence of the ball that needed to be detected, and there was little to attract attention to that.

The next least detected event, the handshakes in the slapping game (31% average), involved hand motions of comparable speed and location to the normal game play, with little break in

action at the transitions. By contrast, the choose-up events in the slapping game (78% detection rate) involved hand motions that covered a larger vertical range, with considerably different finger configurations from the flat views of the hands in the normal game and the handshake events, and that is likely to account for the large increase in detectability.

In their experiment, Becklen and Cervone (1983) found that the umbrella woman was detected at an anomalously high rate during a take in which the ball bounced near her foot as she walked through the ball game. Apparently, that looked like she was kicking the ball, and she was thus included in the perceptual context of the game. In our case, the umbrella woman was detected in 10 of 12 showings of a take in which the hand game players adjusted the position of their hands from the lower portion of the screen towards the middle, just as she was centre screen. This apparently had the effect of motioning toward her, with the hands stopping just below her face. Conversely, in the hand game take that resulted in the fewest detections of the umbrella woman (just 1 detection in 7 showings), the players lowered their hands just as the umbrella woman reached centre screen.

At 75%, the juggler in the ball game was detected at essentially the same high rate as the choose-up event's 78%, and the reasons may be similar. The juggler introduced vertical motions that differed significantly from the regular game play. He also paused right at the centre of the action, creating an uncharacteristic horizontal stability and darkness behind the slapping hands instead of the light backdrop of the ball court normally seen there. If that contrast made a difference, it should yield more detections in the full-colour views of the ball game than the cartooned views, since solid areas are transparent in the cartooned views. Indeed, that does seem to be the case, although the sample sizes are too small to be sure that the effect is real; the juggler was detected 10 of the 12 times he was shown in full colour with a full-colour hand game, and 10 of the 12 times he was shown in full colour with a cartooned hand game. But he was only detected 7 of the 12 times he was shown cartooned against the full-colour hand game.

The slapping game's ball-toss event was the most detected, at 86%, for reasons, we believe, quite different from those above. The slappers' hands and arms disappeared from view briefly when stopping to pick up the ball. It was likely that lack of distracting action (i.e., "noise") that was being sensed, and that drew attention to the hand game just as the ball tossing commenced. It is also possible that the ball itself, being similar in visual extent to the basketball, became contextually relevant to the ballgame task. If so, we would expect a higher detection rate when both views or no views were in full colour, and less when the treatment was mixed, but no such result occurred. But again, the samples are inconclusive, since performance was near ceiling, and the toss was detected in 10 or 11 out of the 12 showings of each treatment.

Although we tried to reproduce N&B's study fairly accurately, our subjects detected the UEs at a much higher rate, clearly detecting 57% of the UEs, while at best 21% were detected by N&B's subjects. Half of their subjects did not detect any of the events, none noted all, and most detections were fairly tentative and incomplete. All but one of our subjects detected at least one UE, and there was little ambiguity in their responses. This has little relevance to the purpose and conclusions of our study, but it does cause us to speculate why. N&B had just 20 attended events during each 1-minute trial, while we had 30 – ostensibly distracting attention more, not less. Due to the nature of the techniques used to combine videos, N&B's subjects had their head motion constrained by a chin rest, while ours did not, which may have made detection easier. Their displays subtended a horizontal visual angle of about 12°, while our's subtended about 19°. The more-compact displays would make it easier to take in more at a time, but would also make features smaller, potentially requiring greater effort to follow. Our videos were of higher technical quality, and used colour, perhaps making the attended tasks easier and the UEs more noticeable. Some of the difference in detection rates can be attributed to differences in the UEs themselves, with the juggler and umbrella woman more detectable

than any of N&B 4 UEs. But where the UEs were intended to match those of the N&B study, our detection rates were still about 3 times more frequent. It may be that in the generation that has passed between the studies, the young subjects themselves have become more attuned to what, for them, is essentially a computer game.

Experiment 2: Detection of same-scene events

In previous studies of IB using overlaid natural scenes, the UEs were in the unattended scene (or an entirely different scene), not in the attended scene. They could not disambiguate two possible causes of IB: inattention to the scene vs. the confusion associated with watching overlaid scenes. In Experiment 2 we conducted a simple test to resolve that ambiguity.

In two of the session sequences ("scripts") from Experiment 1, we simply changed the instructions so that the subjects would attend to the events in two of the UE-containing scenes. Since the presentations were otherwise visually identical to those used in Experiment 1, an increase in detection rate would indicate that it was not the superposition that caused IB.

Methods

Two session scripts from the first experiment were selected that were essentially identical, except that one had the umbrella woman event in trial 8 and the juggler event in trial 12, and the other had the juggler in trial 8 and the umbrella woman in trial 12. The instructions for trials 8 and 12 were modified so that the subject would attend to the scenes that included the juggler and umbrella woman. In the remaining experimental trials, when UEs were presented they were always in the unattended scene, as in Experiment 1. As a further control, in case we did not find a significant difference in detection rates in the two experiments, the UE scene of trial 12 was shown without an overlaying scene, just to confirm that the UEs were easily detectable and would always be reported by the subjects. Otherwise we would have left open the possibility that only the inherent difficulty in detecting a particular UE mattered rather than overlaying.

The juggler and umbrella woman scenes were selected because they were the only scenes with UEs that were not an integral part of the attended action. Hand game UEs were not used in attended scenes since, unlike the ball game events, they were interruptions in the game play rather than incidental to it, and thus would certainly be noticed. The ball game's lost ball event was similarly not suitable for attended viewing. The two scenes used were representative of the events in Experiment 1 noticed most frequently (the juggler) and at a fairly average rate (the umbrella woman). Because we were concerned about the priming effects that a blatantly obvious UE might cause, trial 8 was chosen as the first to have an attended UE, so that we would have some basis for comparing the Experiment 1 and 2 populations before the first attended UE. Either the (unattended) choose-up or handshake event was shown in trial 5, and in trial 6 the hand game was attended while the overlaid ball game had no UEs. The lost ball event was always shown in trial 7. Similarly, since trials after one with a non-overlaid UE would likely prove nothing, trial 12 was used to show the non-overlaid UE. Since we found no significant effect of cartooning on IB in Experiment 1, we simplified matters by always showing the attended UE scene in full colour, and the overlaying scene, if there was one, was cartooned, as that is the combination most representative of the augmented vision devices under development.

15 subjects (age 18–35, 5 male) were tested, alternating use of just the two modified scripts.

Results

Unexpected event detection

All 15 subjects detected all of the UEs that were shown in an attended scene. Given that the average rate from Experiment 1 of detecting the juggler was 0.75, the likelihood that any 15 out of 15 subjects in Experiment 1 would detect the juggler was less than 0.031. Similarly, given the probability of 0.53 that a subject in Experiment 1 would notice the umbrella woman, 15 subjects would all detect her with probability less than 0.0006 (by Fisher exact test for 2×2 tables). Thus the detection rates of 15 out of 15 in Experiment 2 are a significant indication that attending to the scenes with unexpected events caused the difference.

Since all subjects in Experiment 2 detected the UEs when attending the UE scene with another scene overlaying it, it is not surprising that they would also detect the UEs when no overlay interfered. Thus the UE trials without overlays yielded no information about any effect overlaying might have had on detections.

Population equivalence

It is reasonable to question if the Experiment 1 and Experiment 2 population samples were indeed equivalent, as they necessarily (for naïveté) involved a different sample of subjects. Comparing performance with the only 2 subjects in Experiment 1 who were tested with the scripts used in this experiment would not be significant. Rather, we compared detection rates of the handshake, choose-up and lost ball UEs in the two experiments, including only those trials that occurred at or before the first UE detection in a session. This provided larger samples while avoiding priming effects.

Under those conditions, the lost ball event was detected once in Experiment 1 and was not detected in Experiment 2, rendering statistical comparison meaningless.

For the other two events, by comparison of unpaired proportions (Table 3), the two samples are statistically unlikely to have been from different populations ($p > 0.25$).

Experiment 2 Discussion

UEs in an attended scene were always detected. Since videos used in Experiment 2 were identical to ones used in Experiment 1, any degree of confusion caused by the overlapped scenes would have been identical. The only difference (with the minor exception of the scene used as a lead-in to establish which scene was to be attended and to provide timing marks) was that, in the trials of interest, the attended, rather than the unattended, scenes contained the UEs. Thus it was not the overlap, per se, that caused IB, it must have been the relationship of the UEs to the attended scene. This is consistent with the speculated causes in Experiment 1 for some UEs to be noticed much more frequently than others.

General Discussion

In Experiment 1, we examined the effect of cartoon-like edge-filtering on IB. The major, and surprising, result of that experiment is that cartooning had no significant effect on UE detection. This has both positive and negative implications for our augmented-vision devices. We had hoped that cartooning would mitigate the effects of IB, and thus make it easier for UEs to be noticed, such as a child running into the path of a person with peripheral visual field loss, or an automobile accelerating as that person crossed a street. On the other hand, it is fortunate that the cartooning did not exacerbate IB. We know anecdotally (but have not formally tested) that cartooning the minified wide-angle view overlaying a see-through natural view makes it easier to distinguish the views, so that attention can readily be paid to one view or the other. In addition, the cartooned view emphasizes the salient features of the scene, and thus aids

orientation and navigation, making it easier, for instance, to find a door in a corridor, even though (as is so often the case) door and walls are painted the same bland colour. The slight impact that cartooning had on response time to attended events, we believe, is inconsequential compared to the potential benefits. Nonetheless, consciously alternating attention between views requires effort, and it can easily fall prey to the problems of change blindness, so finding ways to mitigate these effects will continue to be a theme of ongoing research.

In Experiment 2, we altered the Experiment 1 protocol slightly to rule out the possibility that it was simply the nature of overlaid presentations – rather than inattention – that caused IB in experiments of this type. When the scene with the UEs was attended, the events were always noticed, even though the presentations were identical. Thus it is not any confusion due to overlaying that causes IB; rather, it is likely the lack of contextual relationship between the attended task and the UE that is the root cause. This, too, bodes well for the use of vision multiplexing in our low-vision aids, as overlaying is not the cause of IB.

The augmented-vision devices we are developing do not show unrelated scenes; they show the same scene at two different scales. It is not clear if that contextual relationship will have a mitigating effect on IB. We are developing an experiment to evaluate that possibility.

Conclusions

Cartoon-like edge filtering had no significant effect on inattentional blindness. Since it is nonetheless likely to be useful in managing divided attention, as well as a way to highlight salient features, it remains a promising technique for our augmented vision aids. We are thus navigating between the Scylla of inattentional blindness and the Charybdis of change blindness, so we continue to seek ways to mitigate their strength. We also conclude that it is not the overlapping of scenes that causes inattentional blindness; rather, it is the contextual separation of the unexpected events from the attended scene.

References

Becklen R, Cervone D. Selective looking and the noticing of unexpected events. Mem Cogn 1983;11:601–608.

Bowers AR, Luo G, Rensing NM, Peli E. Evaluation of a prototype minified augmented-view device for patients with impaired night vision. Ophthal Physiol Opt 2004;24:296–312.

Cartwright-Finch U, Lavie N. The role of perceptual load in inattentional blindness. Cognition 2007;102:321–340. [PubMed: 16480973]

Cohen, J. Statistical Power Analysis for the Behavioral Sciences. Lawrence Erlbaum Associates; Mahwah, NJ: 1988.

Haines, RF. A Breakdown in Simultaneous Information Processing. In: Obrecht, G.; Stark, LW., editors. Presbyopia Research: From Molecular Biology to Visual Adaptation. Plenum Press; New York: 1991. p. 171-175.

Keppel, G.; Wickens, TD. Design and Analysis, A Researcher's Handbook. Pearson Prentice Hall; Upper Saddle River, NJ: 2004.

Koivisto M, Hyona J, Revonsuo A. The effects of eye movements, spatial attention, and stimulus features on inattentional blindness. Vision Res 2004;44:3211–3221. [PubMed: 15482807]

Lavie N. The role of perceptual load in visual awareness. Brain Research 2006;1080:91–100. [PubMed: 16413514]

Lavie N, Hirst A, de Fockert JW, Viding E. Load theory of selective attention and cognitive control. J Exp Psychol Gen 2004;133:339–354. [PubMed: 15355143]

Luo G, Peli E. Use of an augmented-vision device for visual search by patients with tunnel vision. Invest Ophthalmol Vis Sci 2006;47:4152–4159. [PubMed: 16936136]

Mack, A.; Rock, I. Inattentional Blindness. MIT Press; Cambridge, MA: 1998.

Martin-Emerson R, Wickens CD. Superimposition, symbology, visual attention, and the head-up display. Human Factors 1997;39:581–601. [PubMed: 11536851]

Most SB, Simons DJ, Scholl BJ, Chabris CF. Sustained inattentional blindness: The role of location in the detection of unexpected dynamic events. Psyche 2000;6article 14

Most SB, Simons DJ, Scholl BJ, Jimenez R, Clifford ER, Chabris CF. How not to be seen: The contribution of similarity and selective ignoring to sustained inattentional blindness. Psychol Sci 2001;12:9–17. [PubMed: 11294235]

Neisser, U. The control of information pickup in selective looking. In: Pick, AD., editor. Perception and its development: a tribute to Eleanor J. Gibson. Lawrence Erlbaum Associates; Hillsdale, NJ: 1979. p. 201-219.

Neisser U, Becklen R. Selective looking: Attending to visually specified events. Cognit Psychol 1975;7:480–494.

Newcombe RG. Interval estimation for the difference between independent proportions: Comparison of eleven methods. Stat Med 1998;17:873–890. [PubMed: 9595617]

Peli E. Vision multiplexing: an engineering approach to vision rehabilitation device development. Optom Vis Sci 2001;78:304–315. [PubMed: 11384008]

Peli E. Feature detection algorithm based on a visual system model. Proceedings of the IEEE 2002;90:78–93.

Peli, E. Vision multiplexing: An optical engineering concept for low-vision aids (In press). SPIE Proceedings of the conference Current Developments in Lens Design and Optical Engineering VIII; SPIE, Bellingham, WA. 2007.

Peli E, Kim J, Yitzhaky Y, Goldstein RB, Woods RL. Wideband enhancement of television images for people with visual impairment. J Opt Soc Am A 2004;21:937–950.

Peli E, Luo G, Bowers A, Rensing N. Applications of augmented vision head-mounted systems in vision rehabilitation. J Soc Inf Disp 2007;15:1037–1045. [PubMed: 18172511]

Rensink RA. When good observers go bad: Change blindness, inattentional blindness, and visual experience. Psyche 2000;6article 9

Rensink RA. Robust inattentional blindness (abstract). J Vis 2005;5:790.

Simons DJ, Chabris CF. Gorillas in our midst: sustained inattentional blindness for dynamic events. Perception 1999;28:1059–1074. [PubMed: 10694957]

Simons DJ, Levin DT. Change blindness. Trends Cogn Sci 1997;1:261–267.

Stevens, A.; Breslin, C.; Jonas, R.; Mulvanny, P. SID International Symposium Digest of Technical Papers. 29. Society for Information Display; New York, NY: 1998. Evaluation of flat-panel-display technologies for an automotive night-vision HUD system; p. 329-332.

Todd JJ, Fougnie D, Marois R. Visual short-term memory load suppresses temporo-parietal junction activity and induces inattentional blindness. Psychol Sci 2005;16:965–972. [PubMed: 16313661]

Vargas-Martin F, Peli E. Augmented-view for restricted visual field: multiple device implementations. Optom Vis Sci 2002;79:715–723. [PubMed: 12462540]

Appendix A: Balancing Details

This appendix describes in detail the balancing scheme used in Experiment 1.

- 36 subjects were tested.

- Each subject was shown 8 presentations during the IB detection portion of the experimental session, showing all 6 UEs plus 2 trials with no UE (one per game).

- The UE order for each subject was established by a row from a 6×6 digram-balanced square (Keppel and Wickens, 2004), with 6 different squares used to provide the 36 rows. Digram balancing within a square ensured that each UE would be shown immediately before and immediately after each other UE, and using 6 different squares minimized the occurrence of other patterns of repetition.

- The squares were further selected to ensure that the 3×6 subset of each square containing one game's UES contained all 6 possible permutations of the 3 UEs.

- One of the 6 possible permutations of 3 treatments was applied to all of that games UEs in that square, with all 6 used for the 6 squares and a different ordering used within a square for the UEs of the two different games.

- The 4th treatment (Cartoon/Cartoon) was applied to the non-UE trial of each game.

- The order of the 24 possible permutations of 4 attended (non-UE) takes per game was randomized and one permutation was used for each of the first 24 sessions. The first 8 permutations were reused for the last 8 sessions. The unattended non-UE presentation of each game used the same take as the first attended use.

- 36 of a randomized ordering of the 56 possible combinations of 8 things taken two at a time were used to identify the position of the non-UE trials among the 8 presentations. A randomization was selected that ensured that over the full set of 36 sessions, a non-UE presentation was show in each trial the same number of times (9). (It did not, however, ensure that each game was attended in exactly 4 or 5 of the trials at each trial position.).

Appendix B: Trial Order

Trials 1 through 4 (Table B1) familiarised the subject with the games, cartooning, overlaying, and the attention tasks. Trials 5 through 12 (Table B2) are the eight trials in which hit accuracy and response time are scored. Six of those trials included UEs, and provided the basis for the detection rate scores. The two trials that did not include UEs were shown with both views cartooned. One UE trial for each of the two games was shown with both views in full colour, one was shown with just the unattended view cartooned, and the other with just the attended view cartooned. The attended game, cartooning treatment, and UE order of the 8 trials 5 through 12 were balanced between subjects as described in Appendix A.

After Trial 12, the experimenter attempted to determine if the subject had actually detected some events but failed to mention that during the trial questioning. The experimenter explicitly mentioned each UE that the subject had not acknowledged, and asked if the subject had seen that sort of event. If so, the event was re-scored as detected. Descriptions of 3 UE scenes the subject had not been shown (the player substitution, umbrella skip, and one which we had not shown nor even taped: a gorilla) were included in this questioning to catch any tendency for false positives that overly-agreeable subjects might choose, but none was encountered.

Optional trials 13–22 and 24–25 then replayed the UEs the subject had not reported in trials 5–12 or the trial 12 follow-up questioning, as summarised in Table B3. First (potential trials 13–18), each unacknowledged UE-containing presentation was shown exactly as before, but without the distractor task, and the subject was asked the usual post-trial questions. If the subject then noticed a UE, the experimenter asked if it had been noticed in earlier trials, and if so, scored the UE as detected. Then (potential trials 19–22 and 24–25), if still necessary, just the brief UE segment, in full colour, and with no overlaid scene, was shown for each as-yet-

unacknowledged UE. In addition, two UEs were shown, in trials 23 and 26, that had not been shown earlier (but had been mentioned after trial 12; the substitution and umbrella skip), to catch any subject who might always say that a UE had been seen. No such behaviour was detected or suspected.

Table B1

Trials 1–4; practice. The ball game was shown as game 1 for half of the 36 subjects.

	Attended view		Unattended view			Purpose
Trial	Game	Edge?	Game	Edge?	UE?	
1	1	No	None	No	No	Familiarisation with game and distractor task
2	2	No	None	No	No	Familiarisation with game and distractor task
3	1	Yes	None	No	No	Familiarisation with filtering
4	2	No	1	No	No	Familiarisation with overlaid presentation

Table B2

Trials 5–12, the scored trials.. Each subject saw these trials in a different order, with content balanced as described in Appendix A.

	Attended view		Unattended view			Purpose
Trial	Game	Edge?	Game	Edge?	UE?	
varies	1	Yes	2	Yes	No	Blind the experimenter to UE trials.
varies	2	Yes	1	Yes	No	Blind the experimenter to UE trials.
varies	1	No	2	No	Yes	UE trial
varies	1	No	2	Yes	Yes	UE trial
varies	1	Yes	2	No	Yes	UE trial
varies	2	No	1	No	Yes	UE trial
varies	2	No	1	Yes	Yes	UE trial
varies	2	Yes	1	No	Yes	UE trial

Table B3

Trials 13–26. These trials gave the subject opportunities to mention UEs seen in trials 5–12 but not reported earlier. Trials 13–18 repeated, without the distractor task, just the UE trials of 5–12 in which the corresponding UE was not detected nor identified in the post-12 questioning. Trials 19–22 and 24–25 showed just the UE segment of each UE that was still not detected. All sessions included catch trials 23 and 26. The UEs in trials 23 and 26 were not shown earlier, but were mentioned in post-trial-12 questions.

Trial	Purpose
13	Repeat of 1st UE trial if UE was not detected, without distractor task
14	Repeat of 2nd UE trial if UE was not detected, without distractor task
15	Repeat of 3rd UE trial if UE was not detected, without distractor task
16	Repeat of 4th UE trial if UE was not detected, without distractor task
17	Repeat of 5th UE trial if UE was not detected, without distractor task
18	Repeat of 6th UE trial if UE was not detected, without distractor task
19	Choose-up UE segment only, full colour, if it still was not detected
20	Umbrella stroll, UE segment only, full colour, if it still was not detected
21	Juggler, UE segment only, full colour, if it still was not detected
22	Ball toss, UE segment only, full colour, if it still was not detected
23	Umbrella skip UE as catch trial
24	Lost ball, UE segment only, full colour, if it still was not detected
25	Handshake, UE segment only, full colour, if it still was not detected
26	Player substitution UE as catch trial
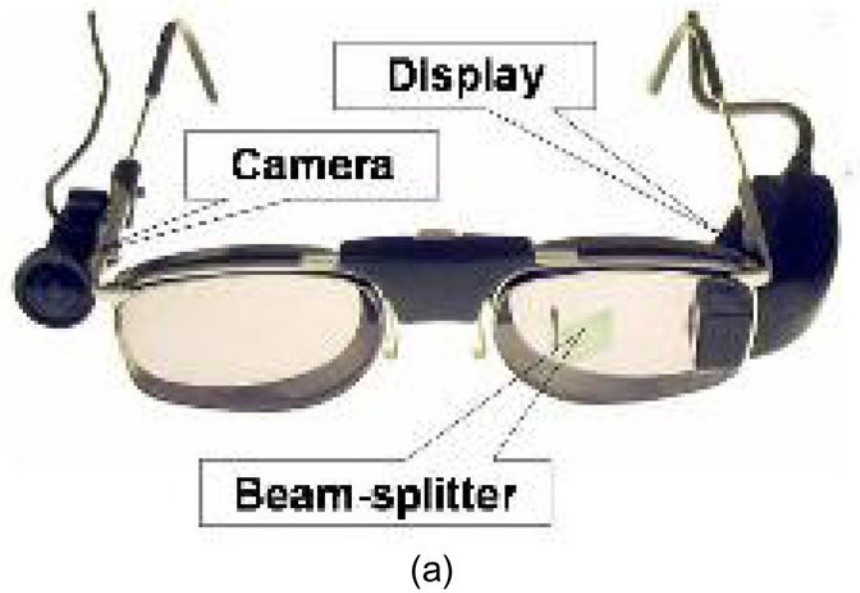
Figure 1.
(a) Augmented-vision spectacles. (b) View through the spectacles. The area outside the VF of a typical user is shaded in this illustration. (c) A detailed view of the central display.
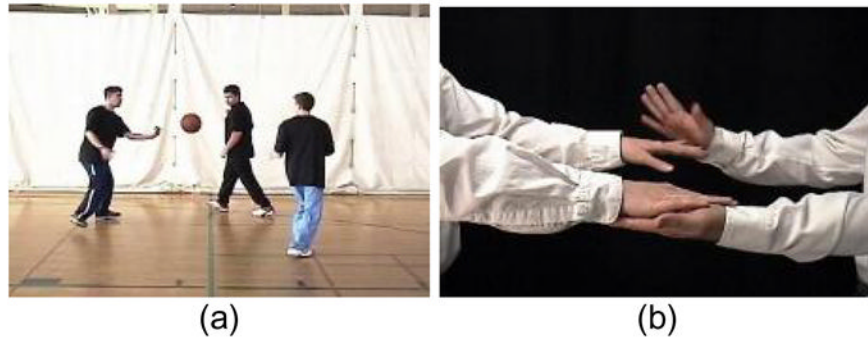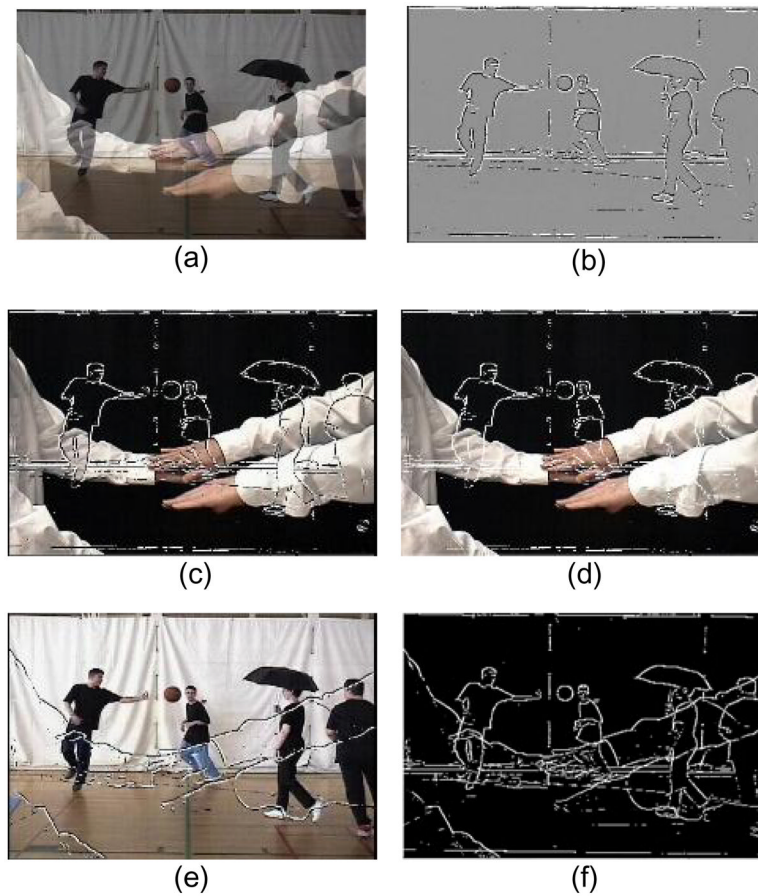
Figure 2.
The ball game (a) and the hand game (b).

SCENE	DESCRIPTION	SECs
(a) Handshake		
	The players stop twice to shake hands once.	2.1 2.3
(b) Choose-up		
	The players stop twice to play three rounds of an odd-even "choose-up" finger game.	7.1 6.0
(c) Ball toss		
	The players stop twice to toss a small ball back and forth twice.	9.6 7.9

Figure 3.
Hand game Unexpected Event (UE) scenes. The on-screen duration of each event episode is given.

SCENE	DESCRIPTION	SECs
(a) Juggler	A man strolls in, pauses at the centre of the game circle, and then strolls off, all the while juggling 3 balls.	5.3
(b) Umbrella stroll	A woman with an open umbrella strolls through the game.	7.1
(c) Lost ball	One player throws the ball out of the game. The players then mime playing without the ball, until the ball is returned.	21
(d) Substitution	A woman replaces one player until that player returns.	20
(e) Umbrella skip	A woman with an open umbrella skips in, stops at the centre of the game circle to jump up and down, and then skips off.	6.3

Figure 4.
Ball game Unexpected Event (UE) scenes. Scenes (d) and (e) were only used as "catch" trials. The on-screen duration of each event episode is given.

Figure 5.
Cartooning. a) Overlaid full-colour views of the ball game and hand game. b) Edge filter output with the ball game view as input. White and black edges appear on a field of grey. c) Cartooned ball game edges over the full-colour hand game. The bipolar edges stand out against light and dark backgrounds. d) The same scene with just white edges (not used in this study). Note the difficulty of seeing them against the bright sleeves. e) Hand game edges over the full-colour ball game. f) Both views cartooned. Only the white edges are used for this condition.
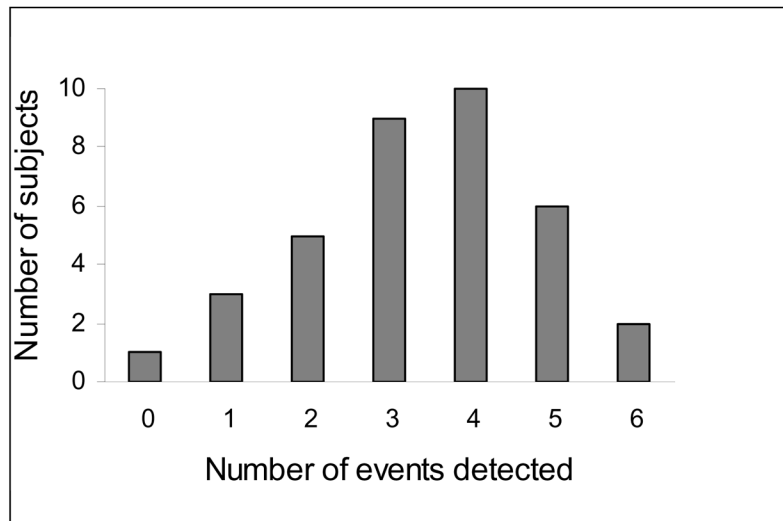
Figure 6.
Frequency of UE detection by subjects. Only 2 out of 36 subjects detected all 6 UEs.
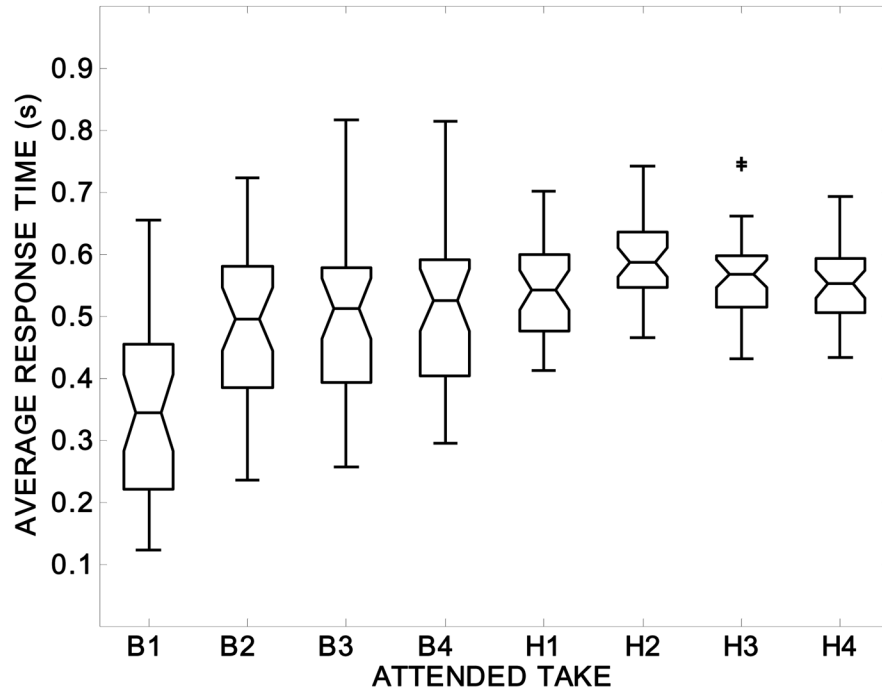
Figure 7.
Average response times to attended events varied across different takes of the attended ball game (B1–4) and hand game (H1–4) scenes. Blocked views during the ball game resulted in wider variations in response time. The boxes enclose the 25th and 75th percentile values, with the median shown at the "waist". The "whiskers" enclose the maximum and minimum values, but H3 excludes two outliers (+) more than 1.5 times the interquartile range above the 75th percentile.

Table 1

Frequency of UE detections by treatment and UE, out of 12 trials for each pairing. (Trials with both views cartooned had no UEs.)

Attended: Unattended:	Full Full	Full Cartoon	Cartoon Full	Total detections
Juggler	10	7	10	27 (75%)
Lost ball	3	2	2	7 (19%)
Umbrella	8	8	3	19 (53%)
Choose-up	9	10	9	28 (78%)
Handshake	4	3	4	11 (31%)
Ball toss	10	10	11	31 (86%)
Total detections	44 61%	40 56%	39 54%	123 57%

Table 2

Effect of cartooning on attended task response time in ms. (±std dev).

| | | Unattended Scene | |
		Full colour	Cartoon
Attended Scene	Cartoon	532 (±84)	522 (±96)
	Full colour	500 (±97)	496 (±100)

Table 3

Population equivalence. The detection rates of equivalent trials in the two experiments were comparable, given the small sample sizes. *p* is the probability that the rates are from equivalent populations.

Detection Rate	Lost Ball	Handshake	Choose-up
Experiment 1	1 of 11 = 0.09	4 of 11 = 0.36	5 of 11 = 0.45
Experiment 2	0 of 8 = 0.0	5 of 8 = 0.63	2 of 7 = 0.29
p	n/a	0.26	0.47