P. Bressan, S. Pizzighello / Cognition 106 (2008) 370–383

The attentional cost of inattentional blindness ☆

Paola Bressan*, Silvia Pizzighello

Dipartimento di Psicologia Generale, Università di Padova, Via Venezia 8, 35131Padova, Italy

Received 9 May 2006; revised 26 February 2007; accepted 3 March 2007

Abstract

When our attention is engaged in a visual task, we can be blind to events which would otherwise not be missed. In three experiments, 97 out of the 165 observers performing a visual attention task failed to notice an unexpected, irrelevant object moving across the display. Surprisingly, this object significantly lowered accuracy in the primary task when, and only when, it failed to reach awareness. We suggest that an unexpected stimulus causes a state of alert that would normally generate an attentional shift; if this response is prevented by an attention-consuming task, a portion of the attentional resources remains allocated to the object. Such a portion is large enough to disturb performance, but not so large that the object can be recognized as task-irrelevant and accordingly ignored. Our findings have one counterintuitive implication: irrelevant stimuli might hamper some types of performance only when perceived subliminally.

Keywords: Inattentional blindness; Subliminal perception; Attention; Awareness; Irrelevant stimuli

1. Introduction

The history of accidents caused by human error is paved with bicycles that crash into runners, cars that crash into bicycles, trains that crash into cars—even planes that land onto other planes (see Haines, 1991, for bloodless evidence). In all these cases the occluding object was there, in full view, but the person responsible for the disaster had

---

☆ This manuscript was accepted under the editorship of Jacques Mehler.

* Corresponding author. Tel: +39 049 827-6691; fax: +39 049 827-6600
E-mail address: paola.bressan@unipd.it (P. Bressan).

doi:10.1016/j.cognition.2007.03.001

not seen it. To detect something, apparently, just looking at it is not enough. When our attention is focused on another object or task, we can fail to see an unexpected object even if it stands exactly where our gaze falls (Koivisto, Hyönä, & Revonsuo, 2004).

Rather appropriately, the significance of the earliest studies on "inattentional blindness" (Neisser & Becklen, 1975; Becklen & Cervone, 1983) went largely unnoticed; namely, these findings were included in textbooks but failed to be incorporated into theories about perception and attention. In these works, people watched two superimposed, semi-transparent videoclips, each representing a simple dynamic scene. Instructed to attend to one, they did not notice an unexpected event happening in the other. Fresh impulse to research came later from the introduction of computer-generated stimuli (Mack & Rock, 1998; Newby & Rock, 1998). Participants were asked to assess which of the arms of a briefly flashed cross was longer. Engaged in the task, about half of the viewers failed to detect the appearance of a small black square in a different location of the screen.

Yet, inattentional blindness managed to attract popular attention only when the role of the unexpected event was taken by a person in a gorilla suit, intruding in a basketball game (Simons & Chabris, 1999). If they were busy counting the number of ball passes between the members of one team, many viewers completely missed the gorilla—even when the beast stopped in the middle of the scene to thump its chest.

Since the introduction of the gorilla, inattentional blindness has rapidly become the focus of high-quality research. An important step in this direction was taken by Most and colleagues (Most et al., 2001; Most, Scholl, Clifford, & Simons, 2005), whose protocol combined a rigorous control of the experimental variables with the use of dynamic (and thus, ecologically more relevant) events as visual stimuli. They devised displays in which a few black and a few white items moved around at random, occasionally bouncing off the edges of the display window. Participants were asked to count the total number of bounces made by either the black or the white items. In these conditions, many failed to detect a new item that unexpectedly moved across the display, even when its shape and color were unique. Remarkably, nearly a third of participants missed a bright red cross traversing a completely achromatic scene.

Taken together, these studies show that we can be unaware of (i.e., unable to report) conspicuous events that take place right in front of our eyes, if our attention is occupied elsewhere. However, the amount of inattentional blindness depends greatly on the personal meaningfulness of the unexpected stimulus (e.g., one's own name is perceived more easily than someone else's, Mack & Rock, 1998), suggesting that even events of which we are unaware must be processed to some degree. Indeed, *total* blindness to the unexpected would not seem very useful biologically. Processing without attention is functional inasmuch as it allows us to maintain our attention on the task at hand, *unless and until* the unattended object is classified as salient enough to become the new focus of attention.

These considerations raise the question of whether the subliminal processing of "unattended" events comes at a measurable attentional cost. On the grounds that what does not enter our phenomenal world cannot harm us, we may assume that whereas seeing the gorilla will affect our counting performance, being blind to the gorilla will not.

This is essentially what previous research has shown (Most et al., 2005): when a new item moves across a dynamic display, accuracy in the main task diminishes significantly for those who notice the intruder and not for those who remain unaware of it. In these experiments, however, participants were asked to keep their eyes on a fixation point, and hence away from both the items they were tracking and any other occurrences of potential interest.

In the work reported here, we compared the effects of superliminal and subliminal “gorillas” by presenting an unexpected, irrelevant item during a visual attention task where participants were free to move their eyes (and therefore, also to inspect directly any event that attracted their attention, as they would do in real life). To our surprise, we found that the unexpected item influenced the task accuracy not when it was seen, but when it was unseen.

2. Experiment 1

2.1. Method

2.1.1. Participants

Sixty participants (21 males and 39 females, mean age 36 years) with normal or corrected-to-normal vision were tested individually. They were randomly assigned to one or the other of two conditions, one visual (single task) and one visual and auditory (dual task). There were 30 participants per condition; one dual-task participant (a nonnoticer) was removed from the analyses because her total number of errors was more than 2 standard deviations above the mean of her condition. The pattern of results did not change if this subject was included.

2.1.2. Stimuli and procedure

The visual stimuli we used were similar to those used by Most et al. (2001), and were presented on a portable computer Toshiba Satellite 1800-412 with a 14” display. On each trial, four black (luminance=1.0 cd/m^2) and four white (luminance=87.4 cd/m^2) L and T shapes moved independently on random paths, at variable velocities, against a 10.6x8.0cm gray (luminance=15.8 cd/m^2) background. Some trials also contained a light gray (luminance=42.3 cd/m^2) cross with the same horizontal and vertical extent as the L’s and T’s, i.e. 8 mm, and the same thickness, i.e. 2 mm. As they moved, the black and white shapes could partially occlude each other, and occasionally bounce off the edges of the display window.

We prepared five separate trials, that were presented in the same order to all participants. The number of bounces was 8 on the first trial, 5 on the second, 6 on the third, 7 on the fourth and fifth. Each trial lasted 12 seconds. Participants were instructed to watch the display and keep a silent tally, using their fingers, of the number of times that the white letters bounced off the edges of the display window. In the dual-task condition, participants also listened to either short stories (comprehension) or lists of words (recall), uttered by a computerized female voice. After each trial, participants reported the number of bounces they had seen; in the dual-task condition, they were also tested on their comprehension or recall.

All observers viewed the display from a distance of about 60 cm and completed five consecutive trials. The first two trials contained no unexpected event. Approximately 2.45 seconds into the third trial (the “critical trial”), the gray cross unexpectedly entered the display from the right side, traversed the screen horizontally along a virtual midline and exited the left side (see Figure 1). The cross remained visible for 7.15 seconds. After this trial, observers answered a questionnaire adapted from Most et al. (2005). They were asked to report whether they had seen anything other than the black and white L’s and T’s, something that was missing in the first two trials. If the answer was yes, they were asked to describe the color, motion direction, and shape of the object. The shape could be picked among four different shapes, graphically represented in the questionnaire: an E-shape, a cross, a heart, and a triangle.

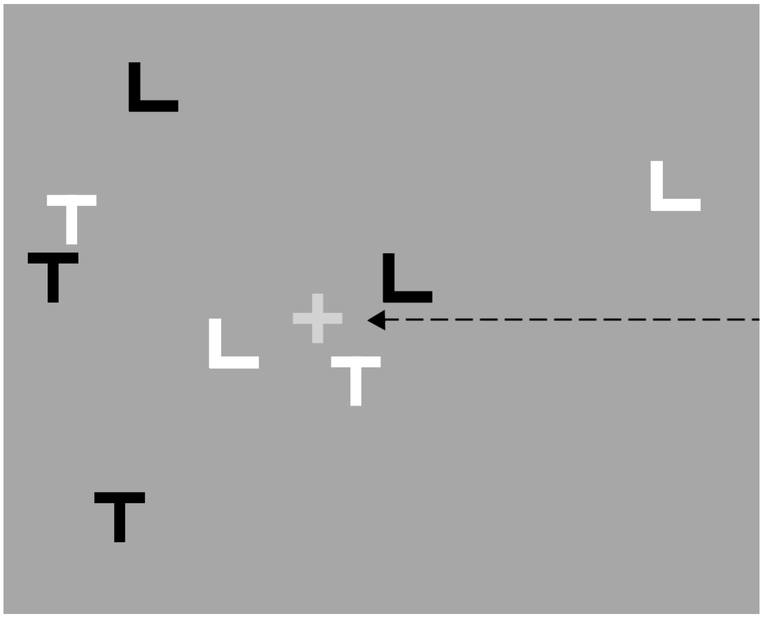

Fig. 1. A single frame of one of the five dynamic displays used in the Experiment. The arrow shows the path traveled by the cross, and was not present in the actual display. The white and black shapes moved around on random paths, occasionally touching and bouncing off the edges of the display window.

Participants then completed a fourth trial (called “divided-attention trial”, because the questionnaire had indirectly alerted them to the possibility that a novel object could appear), after which they answered a second questionnaire, identical to the first. On the fifth and last trial (the “full-attention” trial), participants were simply asked to view the display, without performing any task. After this trial they answered a final questionnaire, identical to the previous two.

2.2. Results

Overall, the cross was noticed by 42% of participants on the critical trial, by 68% on the divided-attention trial, and by 100% on the full-attention trial. The group of 25 observers that mentioned seeing a cross on the critical trial (i.e., reported the shape correctly) includes participants who reported incorrectly the color (10) and/or the direction of motion (3) of the cross; only 9 individuals described all three attributes exactly. A further 3 participants indicated correctly either the color or the direction of motion but not the shape.

Table 1 Mean Number of Errors in the Bounce-Counting Task for Each Experiment

Experiment		Critical trial	Noncritical trial	Critical trial 2
1	Noticers (30)	.57 (.12)	.60 (.15)	
	Nonnoticers (29)	1.31 (.21)	.62 (.17)	
2	Noticers (15)	.47 (.17)	.33 (.16)	.67 (.19)
	Nonnoticers (22)	1.05 (.15)	.32 (.12)	.50 (.14)
3	Noticers (28)	.57 (.13)	.57 (.14)	.50 (.13)
	Nonnoticers (41)	.76 (.13)	.34 (.09)	.61 (.10)
4 (control)	Noticers (30)	.93 (.16)	.43 (.13)	.87 (.15)
	Nonnoticers (8)	1.13 (.40)	1.00 (.27)	1.00 (.27)
5 (control)	Noticers (26)	1.08 (.20)	.88 (.19)	.92 (.18)
	Nonnoticers (13)	.77 (.23)	.85 (.22)	1.23 (.26)
1-2-3 pooled	Noticers (73)	.55 (.07)	.53 (.09)	
	Nonnoticers (92)	.90 (.08)	.42 (.08)	

Note. Mean number of counting errors (standard error of the mean) for noticers and nonnoticers (number of participants), in critical and noncritical trials. Experiments 2-5 contained two critical trials. In the last two cells of the “Critical trial” column, the values are the means of all the critical trials of Experiments 1, 2, and 3 (including those listed in the “Critical trial 2” column).

In this article we will not analyze the comprehension and recall performances and their interferences with the visual task, which will be reported elsewhere (Pizzighello & Bressan, 2006), but focus instead on the effects of inattentional blindness on bounce counting. On the critical trial, participants who had seen a cross made fewer counting errors than people who had not, $t(57)=2.95$, $p=.005$ (all t tests throughout the manuscript are two-tailed). This difference slightly increased, $t(57)=3.13$, $p=.003$, when the group of noticers was enlarged to include all participants who had answered affirmatively to the first question of the questionnaire; that is, those who reported to have seen the unexpected object, whichever the type and number of features correctly described—none for 2 people, at least one for 28 (see Table 1). Hereafter, we will code as "noticers" the participants that answered "yes" to the first question of the questionnaire, and as "nonnoticers" the participants that answered "no". The difference in counting accuracy between noticers and nonnoticers was independently significant in both conditions, as shown in Figure 2.

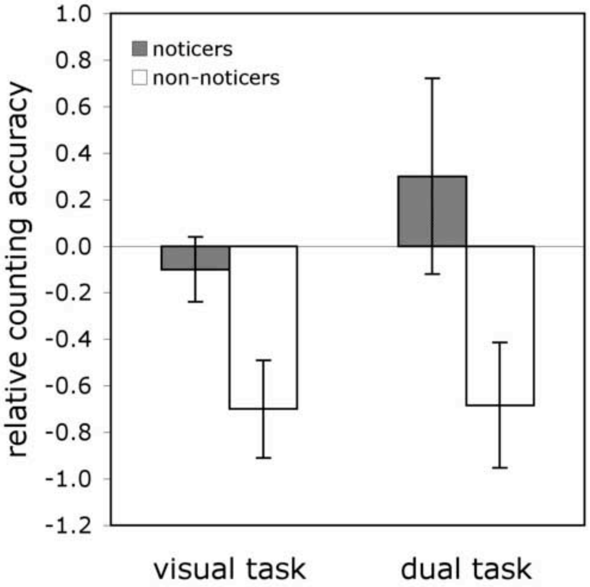

Fig. 2 . Relative counting accuracy on the critical trial, for noticers (i.e., participants who reported seeing the unexpected object) and nonnoticers (i.e., participants who reported not seeing anything). Accuracy is expressed as the absolute deviation from correct number of bounces on the second (pre-critical) trial minus the corresponding number on the third (critical) trial. On the critical trial, participants who saw the unexpected object were on the whole as accurate in counting bounces as they had been on the previous trial, whereas the accuracy of participants who saw nothing decreased significantly, regardless of whether they were only counting (visual task) or also listening to verbal material (dual task). Error bars indicate standard errors.

How can this remarkable result be explained? The possibility that noticers may be finer or more attentive observers in general (and thus, also more proficient in counting the bounces) must be discarded, since in the pre-critical trials noticers were as accurate as nonnoticers, both ts<1; nor was there an accuracy difference between noticers and nonnoticers in the divided-attention trial, $t<1$. A difference emerged only in the critical trial, and was caused by a marked drop in counting accuracy (relative to the pre-critical trial) in nonnoticers, one-sample $t(28)= -3.58$, $p=.001$, but not in noticers, one-sample $t(29)< |1|$. The accuracy drop in nonnoticers consisted in a decrease in the reported number of bounces (they missed on average one bounce more than in the pre-critical trial).

The first criticism that comes to mind is that, after two trials, some participants may have become tired or bored and may thus have paid less attention to the display. This would have made them less likely to notice the new object *and* more likely to make counting errors. After the critical trial a questionnaire had to be filled out, which may have awakened these participants' interest again. This potential scenario could be a trivial explanation of why, unlike noticers, nonnoticers had a drop in performance on Trial 3 (and not on Trials 1, 2, and 4). In Experiment 2, this confound was removed by alternating noncritical and critical trials and presenting the questionnaire only at the end.

3. Experiment 2

3.1. Method

3.1.1. Participants

Thirty-seven participants (18 males and 19 females, mean age 37 years) were tested individually. Three additional participants (1 noticer and 2 nonnoticers) were discarded from the analyses because their total number of errors was more than 2 standard deviations above the mean; the pattern of results did not change if these subjects were included.

3.1.2. Stimuli and procedure

Apparatus, stimuli, and procedure were identical to those of the visual-task condition of Experiment 1, with one important exception. There were four consecutive trials, of which the first and the third contained no unexpected event, whereas the second and the fourth contained the gray cross. Hence, there were no divided-attention and full-attention trials, but a noncritical, practice trial (first), followed by a critical trial (second), followed by a noncritical trial (third), followed by another critical trial (fourth). Observers answered the questionnaire only after the fourth trial. The questionnaire was the same as in Experiment 1; additionally, participants indicated in which trial or trials the unexpected object had appeared. This time, we asked nonnoticers to *guess* the color, motion direction, and shape of the novel object (picking the shape among the four alternatives) even though they claimed they had not seen the object at all.

3.2. Results

Although in this experiment the cross appeared in two separate trials, the proportion of noticers (41%) was essentially the same as in the previous experiment. Interestingly, nonnoticers reported being absolutely sure they had not seen a new object, but when forced to guess picked the correct shape more often (.64) than chance (.25), one-sample $t(21)=3.68$, $p=.001$. A control experiment run on a separate sample of 22 subjects, asked to pick one shape at random from the same set of four shapes, showed that this choice was not due to a bias towards choosing the cross over the alternatives: the cross was not selected (.18) more often than chance (.25), one-sample $t(21)< |1|$.

On the *first* of the two critical trials, participants who had seen a cross made fewer counting errors than people who had not, *t*(35)=2.50, *p*=.017 (see Table 1). The size of the effect was very similar to that found in Experiment 1 (see Figure 2): counting accuracy relative to the noncritical (third) trial was -.73 bounces for nonnoticers, one-sample *t*(21)= - 4.12, *p*<.0001, and -.13 bounces for noticers, one-sample *t*(14)< |1|.

In the second critical trial there was no difference between noticers and nonnoticers, *t*<1, nor relative performance drop for either group. Note that the first and second critical trials are different in one important respect: the cross is a new object on its first passage, but not on the second. For this reason, we ran a new experiment, using two different unexpected objects in the two critical trials.

4. Experiment 3

4.1. Method

4.1.1. Participants

Sixty-nine participants (37 males and 32 females, mean age 29 years) were tested individually. Four additional participants (1 noticer and 3 nonnoticers) were discarded from the analyses because their total number of errors was more than 2 standard deviations above the mean; the pattern of results did not change if these subjects were included.

4.1.2. Stimuli and procedure

Apparatus, stimuli, and procedure were the same as in Experiment 2, except for two details. First, the two unexpected objects were a circle and a diamond rather than two crosses. They had the same color (light gray), the same horizontal and vertical extent, and the same thickness as the cross. The circle and the diamond appeared in the second and fourth trials respectively for half of the participants, and in the opposite order for the other half.

Second, the number of bounces was counterbalanced across trials. For half of the participants the bounces were 7 in the first trial, 6 in the second and third trials, and 7 in the fourth; for the other half, the bounces were 6 in the first trial, 7 in the second and third trials, and 6 in the fourth.

The questionnaire was the same as in Experiment 2, except that shapes could be chosen among eight different alternatives, depicted in random order: a circle, a diamond, a square, a triangle, a cross, an E-shape, an H-shape, and an X-shape.

4.2. Results

Only 28 (41%) of our observers reported seeing the novel object, and correctly recalled either the first of the two shapes (8), or the second (12), or both (8). Nonnoticers were sure they had not seen any unexpected object, but when forced to guess picked either the circle or the diamond more often (.44) than chance (.25), one-sample *t*(40)=2.41, *p*=.021. As in the previous experiments, noticers and nonnoticers performed differently in the critical and noncritical trials, as shown by the significant interaction between trial (critical, noncritical) and group (noticers, nonnoticers), $F(1, 67)= 4.39$, $p=.040$. For nonnoticers, counting accuracy relative to the third (noncritical) trial was lower both in the first critical trial, one-sample $t(40)= -2.80$, $p=.008$, and in the second critical trial, one-sample $t(40)= -2.43$, $p=.020$. For noticers, relative counting accuracy was not significantly different from zero in either the first or the second critical trial, both $ts(27)< |1|$.

It seems rather unlikely that most of our participants, in perfect unison, would withdraw their attention from the display on the second trial, become suddenly attentive on the third, and withdraw attention yet again on the fourth trial. Taken together, Experiments 2 and 3 show that the interference of the unexpected object on performance is a robust one, and it is not due to accidental variations in attention over time. The unexpected, irrelevant object diverts attentional resources from the main task not when it is seen, but when it is unseen. Averaging across the three experiments, failing to notice the unexpected object led to more counting errors than noticing it, independent-samples $t(163)=3.09$, $p=.002$ (see Table 1). Relative to noncritical trials, counting accuracy in critical trials dropped for nonnoticers, $t(91)=5.65$, $p<.0001$, but not for noticers, $t(72)=0.20$, $p=.839$.

One potential, not very exciting, explanation is that attentional resources decrease purely because they are spread thinner in a trivial way. Noticers, by definition, discriminate the deviant item from the items they need to track, classify it as irrelevant to the task and hence manage to ignore it. Nonnoticers, being unable to identify the new object, may simply mistake it for an additional white letter, thereby adding to the five relevant items a further one. This is of course a harder, more error-prone task. Experiment 4 was designed to tackle this issue by actually presenting an extra white letter (a B or an S) in some of the trials. We reasoned that, if the pattern of errors for people who truly have to keep track of an extra white letter is the same as the pattern of errors for the nonnoticers in the previous experiments, the “additional-item” explanation would be the most parsimonious one.

5. Experiment 4

5.1. Method

5.1.1. Participants

Thirty-eight participants (10 males and 28 females, mean age 31 years) were tested individually. Four additional participants (all noticers; the pattern of results did not change if these subjects were included) were removed from the analyses either because their total number of errors was more than 2 standard deviations above the mean (2), or because they spontaneously reported counting bounces for the B and S (2).

5.1.2. Stimuli and procedure

Apparatus, stimuli, and procedure were the same as in Experiment 3, with the only exception that the unexpected object was a white letter (either a B or an S) rather than a gray shape (either a circle or a diamond). The extra letter had the same color, horizontal and vertical extent, and thickness as the white L's and T's. The B and the S appeared in the second and fourth trials respectively for half of the participants, in the opposite order for the other half. Note that, as before, participants were asked to count the number of times that the white letters bounced off the edges of the display window, which automatically made the additional white letter a relevant item, even though it never bounced.

The questionnaire was the same as in Experiment 3, except that among the final eight shapes the circle and the diamond were replaced by a B-shape and an S-shape.

5.2. Results

Thirty participants out of 38 (79%) reported seeing the unexpected white letter, and correctly recalled either the first of the two unexpected letters (3), or the second (7), or both (20). In the questionnaire, the nonnoticers failed to pick either the B or the S (.125) more often than chance (.25), one-sample $t(7)= -1$.

The counting accuracy for the 8 nonnoticers did not decrease in the trials containing the additional white letter, $t(7)<1$; their performance was actually the same across all trials (presumably, these participants remained focused on the L's and T's). The counting accuracy for the 30 noticers worsened in the trials containing the additional white letter, $t(29)= -3.04$, $p=.005$ (see Table 1). The effect was entirely due to a significant increase in the mean reported number of bounces relative to the noncritical trial: on average noticers counted an extra 0.5 bounces, one-sample $t(29)=2.98$, $p=.006$. This figure is significantly higher than the corresponding one for the nonnoticers of Experiments 1, 2, and 3, independent-samples $t(120)= 3.99$, $p<.0001$. Lower performance in our previous experiments consisted mainly in a larger number of bounce misses: indeed, if only participants that missed (as opposed to added) bounces on at least one of the two critical trials are included, the performance drop is significant in Experiments 1, 2, and 3 (nonnoticers, all ts> 3.0, all ps<.01), but not in Experiment 4 (either noticers or nonnoticers, both ts< |1|). The pattern of errors in the two blocks of experiments is therefore clearly different.

These results suggest that it is unlikely that the inaccuracy induced by the unseen object in the previous experiments was caused by mistaking the object for an additional relevant item. In Experiment 4, when the identity of the extra item went unnoticed the counting performance did not change (whereas in the previous experiments it worsened), which implies that the new object had failed to attract either implicit or explicit attention—as also speculatively suggested by the fact that, unlike in the previous experiments, shape guessing was totally random. In order to affect performance the new object had to be processed deeply enough to be included among the relevant items, but as a consequence it was recognized (whereas it remained unidentified in the previous experiments). Moreover, treating the new object as a relevant item increased the mean reported number of bounces, and did not decrease it as in the previous experiments.

A result that appears partly at odds with our main finding comes from the work of Most et al. (2005). These authors report that, in their experiments, the critical-trial counting accuracy diminished for both noticers and nonnoticers, with a more pronounced effect for the noticers. Unlike us, however, Most et al. used a fixation point. Because this could be important for understanding the discrepancy between our results and theirs, we ran a new experiment, identical in all respects to Experiment 3 except for the presence of a fixation point.

6. Experiment 5

6.1. Method

6.1.1. Participants

Thirty-nine participants (10 males and 29 females, mean age 24 years) were tested individually. One additional participant, a noticer, was discarded from the analyses because her total number of errors was more than 2 standard deviations above the mean. The pattern of results did not change if this subjects was included.

6.1.2. Stimuli and procedure

Apparatus, stimuli, and procedure were the same as in Experiment 3, with the only exception that observers were instructed to fixate a small blue square located in the center of the screen, similar in all respects to the fixation point used by Most and colleagues (Most et al., 2001; Most et al., 2005).

6.2. Results

Twenty-six participants out of 39 reported seeing the novel object, and correctly recalled either the first of the two shapes (7), or the second (8), or both (9), or neither (2). Although the only difference between Experiments 5 and 3 was the presence of a fixation point, the results were markedly dissimilar. First, the probability of noticing the unexpected object was significantly larger in Experiment 5 (fixation point: 67%) than in Experiment 3 (no fixation point: 41%), independent-samples $t(106)=2.67$, $p<.009$. This may not be surprising if one considers that the unexpected object, during its translation across the center of the screen, passed right behind the fixation point (like in the experiments of Most and colleagues). Second, unlike in Experiment 3, nonnoticers did *not* pick either the circle or the diamond (.38) more often than chance (.25) in the final questionnaire, one-sample $t(12)<1$. This suggests that any subliminal processing of the novel object was reduced relative to Experiment 3.

Third, and most importantly, for nonnoticers counting accuracy relative to the third (noncritical) trial was *not* significantly different from zero in either the first critical trial, one-sample $t(12)<1$, or the second critical trial, one-sample $t(12)=1.44$, $p=.175$ (see Table 1). For noticers, it was not significantly different from zero in either critical trial, both ts(26)<1.

The divergence between the results of Experiments 3 and 5 suggests that the unseen, irrelevant object draws attentional resources away from the main task when participants are free to track the moving targets, but not (or not nearly as much) when they are forced to fixate a point. A parsimonious explanation stems from the fact that in the latter case the counting task is more difficult. Tracking a number of items in chaotic motion closely enough to detect whether they touch the borders of the screen is much harder if one is unable to look at them directly. Indeed, averaging across the three trials of interest, the number of counting errors was significantly larger in Experiment 5 (.96) than in Experiment 3 (.56), $t(106)=3.79$, $p<.0001$, and this was true for both noticers and nonnoticers (both *p*s<.03). More difficult tasks demand more attention; more attention devoted to primary tasks implies that less attention is available for secondary activities, such as subliminal monitoring of extraneous objects. Processing of task-irrelevant stimuli has indeed been shown to depend upon the level of cognitive load in the relevant task (e.g., Lavie, 2005).

7. General discussion

In our no-fixation experiments, only about 40% of observers noticed the unexpected, irrelevant stimulus, a remarkable finding considering that the display could be explored freely, that in two of three experiments the stimulus itself appeared not on one but on two separate trials, and that each time it was clearly visible for 7 seconds. Our data suggest that, if it entered awareness, the novel object could be rapidly classified as unrelated to the visual task (typically, because of a mismatch in the "shape" dimension) and disregarded, at no extra cost. If it was not consciously perceived, however, the novel irrelevant object lowered performance in the task. This shows that attention capture can occur without awareness (McCormick, 1997; Most et al., 2005; see also Rensink, 2000, 2004), and that this type of attention capture is costly.

The decline in performance in people who did not see the new irrelevant stimulus is hard to explain as a combination of failure-to-identify *and* mislabeling-as-relevant, because, in fully comparable circumstances (Experiment 4), relevant stimuli influenced performance if and only if they were identified as such. Furthermore, the way in which additional relevant items affected counting performance was significantly different from the way in which additional irrelevant items did. Consistent with this argument is the fact that, although letter shapes were always present among the questionnaire alternatives, nonnoticers did not tend to pick those. Indeed, they picked the correct object (cross, circle, or diamond) more often than chance. Such a choice is suggestive of subliminal processing.

In Most et al. (2005), any accuracy decrease for nonnoticers was always quite slight, and became significant only when the data from 7 experiments were pooled. The results of our Experiment 5, where we also failed to find a significant decrement in performance for nonnoticers, suggests that the presence of a fixation point in Most et al.'s experiments may have been the main factor responsible for this data pattern. Arguably, fixation makes the tracking task harder, thereby increasing the cognitive load and leaving less attention available for subliminal monitoring of novel objects (e.g., Lavie, 2005).

Most et al.'s (2001, 2005) counting task was more cognitively demanding than ours in two further respects. First, in our study subjects used their fingers to count the bounces, whereas in Most et al.'s they kept the count entirely in their head. Second, in Most et al.'s critical trials, that lasted 15 seconds, the target letters bounced on average 18 times (range 10-26; S. Most, personal communication, July 2006). In our own critical trials, that lasted 12 seconds, the target letters bounced on average 6.5 times (range 6-7). The high rate of bouncing may economically explain why, unlike us, Most et al. found a sizeable decrease in accuracy for noticers. At an average rate higher than 1 bounce every second, even a relatively short inspection of the novel object was bound to divert the attention of an otherwise perfect noticer from the moving targets long enough that a couple of bounces would be missed. At our average rate of 1 bounce every 2 seconds, that would be less likely the case. It therefore stands to reason that, in Most et al.'s experiments, the noticers' advantage in being able to disregard the irrelevant object was offset by the large probability of missing bounces when attention was diverted from the main task to the object itself. Furthermore, whenever it was gray rather than having either the attended or the to-be-ignored colors, the unexpected object was quite inconspicuous against the gray background (Weber contrast ranged from -0.4 for dark-gray unexpected items to 0.5 for light-gray ones, whereas our cross had a Weber contrast of 1.7), hence adding to the time needed to recognize it and dismiss it (see Burkhardt, Gottesman, & Keenan, 1987).

In the real world, where we are not normally requested to keep fixating a point while things happen around us, attentional shifts tend to be associated with eye movements. Thus, our experimental procedure (which reveals a significant accuracy drop in nonnoticers) may mirror daily-life conditions more closely than Most et al.'s (which does not). Explicit attention capture presumably triggers an eye movement to the object, and prompt recognition of the object's irrelevance to the task will re-direct the eye back to the original target. Implicit attention capture, on the other hand, consumes cognitive resources without rewarding us with the information needed for shifting attention strategically. Outside the laboratory, hence, irrelevant stimuli may hamper some types of performance only when perceived subliminally.

Our data may help to take a stand in one of the thorniest issues about inattentional blindness, still unsettled in the literature. Does the failure to report unexpected items reflect a sort of blindness (a failure to perceive them)—or rather a form of agnosia (a failure to recognize them, see Simons, 2000) or amnesia (a failure to remember them, see Wolfe, 1999)? Our finding of a deterioration in counting accuracy in nonnoticers, but not in noticers, appears inconsistent with the idea that unattended items are either "perceived-but-miscategorized" or "perceived-but-forgotten". Such accounts assume that nonnoticers *do* see the unexpected object, but then proceed either to misclassify it as expected or to forget it; they presuppose that *all* observers are actually noticers. Yet if this were the case, the cost should be equally large, or larger, when the object is assigned to the more distracting "unexpected" category (as noticers do) rather than to the less distracting "no-news" one (as nonnoticers do), or when a trace of the object is kept in memory until the end (as noticers do), rather than disposed of (as nonnoticers do). That is, performance should, if anything, be worse for noticers than for nonnoticers.

In conclusion, the failure to report objects must be a form of blindness; though a peculiar one, to be sure. We propose that an unexpected stimulus falling outside the focus of attention produces a state of alert that, normally, would generate an attentional shift. In conditions in which the diversion of attention is inhibited by an absorbing task, a fraction of the available attentional resources remains allocated to unconscious processing of the new stimulus. If and when the latter surpasses a salience threshold, explicit attention is attracted automatically and conscious perception ensues. Relevant stimuli cross the boundary for (nearly) everyone; irrelevant stimuli cross the boundary for some people and not for others. Superliminal gorillas can be recognized as task-irrelevant and thereby ignored, but subliminal gorillas will not go away. Apparently, inattentional blindness is not *that* inattentional after all.

Acknowledgments

Part of this research (Experiment 1) was submitted as a Master's thesis by Silvia Pizzighello to the Department of General Psychology at the University of Padua, under the supervision of Paola Bressan. We are indebted to Stefano Vezzani for precious discussion, and for first bringing the outstanding work of Most and colleagues to our attention. Thanks to Peter Kramer, Daniel Simons, Steven Most, James Cutting, Carlo Umiltà, and one anonymous reviewer for providing useful comments; additional thanks to Steven Most for sending us requested data.

References

Becklen, R., & Cervone, D. (1983). Selective looking and the noticing of unexpected events. *Memory and Cognition, 11*, 601-608.

Burkhardt, D. A., Gottesman, J., & Keenan, R, M. (1987). Sensory latency and reaction time: dependence on contrast polarity and early linearity in human vision. *Journal of the Optical Society of America, 4,* 530-539.

Haines, R. F. (1991). A breakdown in simultaneous information processing. In G. Obrecht & L. W. Stark (Eds.), *Presbyopia Research: From Molecular Biology to Visual Adaptation* (pp. 171-175). New York: Plenum.

Koivisto, M., Hyönä, J., & Revonsuo, A. (2004). The effects of eye movements, spatial attention, and stimulus features on inattentional blindness. *Vision Research, 44,* 3211-3221.

Lavie, N. (2005). Distracted and confused? Selective attention under load. *Trends in Cognitive Sciences, 9,* 75-82.

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

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

Most, S. B., Scholl, B. J., Clifford, E. R., & Simons, D. J. (2005). What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review, 112,* 217-242.

Most, S.B., Simons, D.J., Scholl, B.J., Jimenez, R., Clifford, E., & Chabris, C.F. (2001). How not to be seen: The contribution of similarity and selective ignoring to sustained inattentional blindness. *Psychological Science, 12,* 9-17.

Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. *Cognitive Psychology, 7,* 480-494.

Newby, E., & Rock, I. (1998). Inattentional blindness as a function of proximity to the focus of attention. *Perception, 27,* 1025-1040.

Pizzighello, S., & Bressan, P. (2006). *Auditory Attention Causes Visual Inattentional Blindness.* Manuscript submitted for publication.

Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition, 7,* 17-42.

Rensink, R. A. (2004). Visual sensing without seeing. *Psychological Science, 15,* 27-32.

Simons, D. J. (2000). Attentional capture and inattentional blindness. *Trends in Cognitive Sciences, 4*, 147-155.

Simons, D.J., & Chabris, C.F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28,* 1059-1074.

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting Memories: Cognition of Brief Visual Stimuli* (pp. 71-94). Cambridge, MA: MIT Press.