The Visible Gorilla: Fast-Moving Objects Override Endogenous Attention

Authors:

Pascal Wallisch
Department of Psychology
New York University
New York City, New York, USA
pascal.wallisch@nyu.edu

Wayne E Mackey
Department of Psychology
New York University
New York City, New York, USA

Michael W Karlovich
New York University
New York City, New York, USA

David J Heeger
Department of Psychology and Center for Neural Science
New York University
New York City, New York, USA

Abstract: It is widely believed that the detection rates of unattended objects are low - even if they are moving. Here, we report the results of three high-powered experiments indicating that these detection rates are strongly modulated by the relative speed of the unattended object. Specifically, fast – but not slow – objects are readily detectable, whether they are attended or not. These results are consistent with a sentinel system that focuses on the task at hand unless overridden by the detection of a strong motion signature in the periphery.

Keywords: Selective attention, endogenous attention, exogenous attention, gorilla, inattentional blindness, stimulus speed, motion perception, mTurk, power

Acknowledgments: We would like to acknowledge Tuvia Lerea for creating the JGL environment to run the experiment online. We would like to acknowledge Dan Grossman for coding participant responses and we would like to acknowledge the NSF for grant DGE 1342536 to WEM.

Introduction

Attention – focusing processing resources on a location, feature or object in the environment – is a central characteristic of cognition (James, 1890). Effects of attention are reliable and strong, manifesting behaviorally as shorter reaction times (Posner, 1980) and neurally in terms of modulated firing rates of single neurons (Desimone, 1998; Treue & Maunsell, 1996), resulting in a boost of signal-to-noise ratio for attended signals (Maunsell, 2015). Consequently, focusing on relevant information is critical for the success of the organism, as dividing attention on several tasks simultaneously has been consistently linked to deleterious effects on performance (Pashler, 1992; Pashler, 1994; Strayer 2006). Accordingly, unattended information remains largely unprocessed and tends to not reach awareness. For instance, auditory material from unattended ears is not processed semantically (Cherry, 1953), unattended visual stimuli are not readily perceived (Rock, Grant & Mack 1992). This “inattentional blindness” extends to the inability to detect unattended changes, commonly known as “change blindness” (Simons & Levin, 1997; Simons, 2000; Simons, 2005). This is surprising, as the visual system generally emphasizes changes of luminance over space (Barlow, 1953; Cugell & Robson, 1966) and time (Müller & Lennie, 1999; Holcombe, 2009). The combination, changes of spatial position over time – or visual motion – is especially salient, a phenomenon known as attentional capture (Hillstrom & Yantis, 1994). Whereas it is debated why this is the case – whether it is the appearance of a new object, or motion per se that triggers attentional capture, motion onsets are considered to be particularly effective in capturing attention (Abrams & Christ, 2003), as reflected by rapidly adapting neural responses to motion signals (Priebe, Churchland & Lisberger, 2002). So one could expect moving objects to be readily detected, as they enter – and traverse – a scene.

However, it has been reported that even moving stimuli will remain largely undetected, and do not trigger the redeployment of attention, when engaged in a demanding task (Simons & Chabris, 1999). This view – that even dynamic unexpected stimuli do not override the deployment of task-focused attention has become increasingly influential (Galotti et al., 2013; Goldstein & Brockmole, 2016), to the point where it has been suggested that it is mere exposure time – not speed – that determines the detectability of moving objects (Kreitz et al., 2016).

Here, we show that this widely accepted view is false. We found that slowly moving stimuli fail to capture attention, replicating previously reported results. But we also found that fast stimuli do capture attention, consistent with the operation of a speed-based sentinel system – allowing an organism to focus on the task at hand unless overridden by the detection of a strong motion signature of an unexpected stimulus.

Experiment 1

In our first experiment, we aimed to test the hypothesis posited by Simons & Chabris (1999) that dynamic stimuli fail to capture attention. We were interested in whether changing the speed of the unexpected stimulus - in this case a gorilla - would make it more detectable.

Method

Participants

We recruited 1624 participants on the Amazon Mechanical Turk platform and randomly assigned them to experimental conditions. We excluded participants that reported to be already familiar with the task, reported low video quality or other technical problems during stimulus presentation, reported having ADD or taking Adderall or reported that they were not doing the task. In addition, we excluded participants that were more than one pass off in their reported pass counts in the counting task. This last criterion might seem overly strict, but this task is rather easy if someone is paying attention, so it could be argued that even someone who is off by 2 passes was not sufficiently focusing on the task, as instructed. We retained data from 929 participants for further analysis. None of our conclusions depend on whether we included all participants or adopt even more stringent criteria (e.g., only keeping data from participants with perfect pass counts). We adopted this set of criteria so as to strike a reasonable balance between keeping only people who were credibly engaged in the task and without obvious confounds (such as expecting the gorilla) while still retaining the bulk of our dataset, from data logged in an online environment. All procedures were approved by the University Committee on Activities Involving Human Subjects (UCAIHS) at NYU.

Stimuli and Task

We used 4 videos in this task. One was the original “gorilla” video used by Simons & Chabris (1999). In this video, two teams of 3 “basketball players”, one wearing white shirts (“the white team”) and one wearing black shirts (“the black team”) move about in front of elevators and pass basketballs to each other. At some time around the middle of the video, someone wearing a gorilla suit enters the screen from the right then stops in the middle before leaving the screen on the left, 10 seconds after entering. In the other three videos, we replicated this setup in front of elevators at NYU (see supplementary materials for the videos). In two of those videos, the gorilla takes 6 and 2 seconds to traverse the scene. In the remaining video, the gorilla takes 10 seconds to traverse the screen - as in the original video - but instead of walking, the gorilla moves in a more biologically appropriate fashion, i.e., in leaps. Participants were instructed to count either the number of passes made by the “black” or the “white” team, respectively. So there were 8 experimental conditions (4 videos x 2 conditions), to which participants were assigned randomly. Importantly, each participant did only a single trial, as this task is only valid the first time they do it - after one trial, participants would have an expectation that something unexpected might happen, which transforms the task to a search task. This is why this task is especially suited for online data recording via Mechanical Turk - it is short and allows us to recruit a large number of participants. The videos themselves were hosted on Youtube, so as to ensure smoothness, but with disabled video controls (i.e. participants could not rewind or pause the video).

Measures and Analysis

After presenting the video, we asked participants how many passes they counted. After that, we asked whether they noticed anything unexpected and if so, what. Responses to the latter question were independently coded by two raters. These raters were blind as to the experimental conditions that the participants were assigned to. Raters judged whether the response of the participant indicated that the participant detected the gorilla or not (the inter-rater agreement was 0.99). We also asked them whether they already knew this task, whether they participated in a study like this before, whether they have ADD or are taking Adderall and whether they were actually doing the task.

As the key dependent variable in this task is categorical (reporting seeing something unexpected or not), we use Chi square analyses throughout this paper. All data were analyzed using MATLAB (Mathworks, Inc.).

Results

Was there an overall effect of paying attention to passes made by the white vs. black team? Participants who were instructed to count the passes made by the black team were much more likely to detect the unexpected gorilla that entered the scene than those counting the passes made by the white team, replicating Simons & Chabris, (1999). There was evidence of a strong effect of deploying endogenous selective attention to passes by the black vs. white team $\chi^2 (1, N = 929) = 68.33, p < .001$; Cramer's $V = 0.271$.

Did gorilla speed influence the probability that the unexpected event was detected? This was the case, for those who were paying attention to passes made by the white team (Figure 1a): $\chi^2 (2, N = 379) = 10.52, p = .005$; Cramer's $V = 0.167$. But there was no evidence for a difference between speed conditions when participants were paying attention to the passes of the black team: $\chi^2 (2, N = 310) = 0.63, p > .250$.

-- FIGURE 1 about here --

Figure 1: Proportion of participants reporting a relevant unexpected event. a) x-Axis: Time the gorilla spent traversing the screen, in seconds. y-Axis: Proportion reported. Solid line: Participants reporting passes by the Black team. Dashed line: Participants reporting passes by the White team. Error bars indicate standard error of the mean. b) As in a), but comparing the performance of participants in the smoothly walking vs. leaping gorilla (otherwise time matched - spending 10 s on the screen) condition.

These results indicate that the speed of unexpected objects matters in terms of their detection probability, but this is only apparent for the participants paying attention to the white passes, as the performance of those paying attention to the passes of the black team are already close to the ceiling for slow speeds. but only if the unexpected object does not share the most salient distinguishing feature (as per instruction here its black color) that participants are already paying attention to.

Was there an effect on detection rate of how the gorilla moved, regardless of translational speed? In the original video (Simons & Chabris, 1999), the gorilla moved slowly in terms of translational motion but also walked upright like a person and smoothly. We were wondering whether a gorilla that spends the same overall time on the screen as in the original video (10 seconds), but has sharp local motion transients due to leaps would be more detectable. We found that this was the case for participants that reported counts by the white team (Figure 1b): $\chi^2 (1, N = 265) = 13.27, p < .001$; Cramer's $V = 0.22$. Again, there was no evidence for a difference between the two conditions when participants counted the passes made by the black team: $\chi^2 (1, N = 175) = 0.62, p = .43$.

Discussion

We replicated the main result of Simons & Chabris (1999) - that participants are surprisingly likely to miss unexpected events, if they are unattended, even if they are in plain view. However, we also found that object speed strongly modulates detectability by participants who are not already attending to the stimulus feature of the unexpected object. We would like to emphasize how remarkable this finding is in light of Kreitz et al. (2016). In the fast moving gorilla condition, participants had only a fifth of the time to detect the presence of the gorilla compared to the slow condition, yet they reported seeing it much more readily. This suggests that it is not merely exposure duration, but rather object speed that triggers an override of attention. This is highly counterintuitive: The less long the gorilla is on the screen, the more likely participants were to see it. In addition, it is not only global object speed that matters - when we matched the exposure duration to that of the original gorilla video, but have the gorilla move by fast jumps instead of walking smoothly, it was also much more detectable. This strongly suggests that it is fast speeds - be they local or global - that drive the override in attention.

However, this experiment had several limitations. First, it is difficult to do a “gorilla” style experiment almost 20 years after Simons & Chabris. Presumably - due to the overwhelming popularity of the original results - it is now hard to recruit truly naive participants. The original paper reported detection rates of 8% (for the white team condition) and 67% (for the black team condition), in contrast to around 60% (for the white team condition) and over 90% (for the black team condition) when we reused the exact same video. For whatever reason, our participants saw the unexpected relevant stimulus (the gorilla) much more readily. This could also be due to the fact that we recruited participants in an online environment or other factors that we cannot distinguish.

Second, the gorilla-detection task is difficult to control and parameterize. There are all kinds of differences between videos in natural settings other than the speed of the gorilla or the color of the shirt of the players that might make a difference, e.g. number of passes in the video, passing patterns or any other number of subtle differences that are hard to match between conditions. However, it is encouraging to both replicate the original effect and consistently extend it to new speed conditions despite these limitations and despite logging the data online, which is a notoriously noisy environment. The fact that we were able to do this speaks to the robustness of the effect. We would be concerned if the effects relied on minute details of the presentation or other boundary conditions. Yet, it is important to be able to characterize the effect of speed on the override of endogenous attention in a more controlled and parametric fashion and with a task that is not already so familiar to so many participants - but that is conceptually similar. This is what we attempted to do in experiment 2.

Experiment 2

To overcome some of the limitations encountered when doing a version of the original gorilla task - like experiment 1 - we created a novel, but conceptually similar task to probe the ability of participants to detect unexpected relevant stimuli.

Method

Participants

Participants were recruited through the Amazon “Mechanical Turk” platform. We recorded data from 1430 participants. We excluded participants using similar criteria as in experiment 1, retaining data from 1185 participants for further analysis. Participants were randomly assigned to one of 24 experimental conditions. As in experiment 1, all procedures were approved by the University Committee on Activities Involving Human Subjects (UCAIHS) at NYU.

Stimulus and Task

Before starting the task, we asked participants to provide screen parameters (specifically screen size and aspect ratio) to ensure that the stimulus appeared in a similar fashion on a wide variety of screens. In addition, we asked participants to maintain a viewing distance of 2 feet from the screen, consistent with the preferred (Jaschinski, 2002) and recommended (Rempel et al., 2007) viewing distance. Assuming this viewing distance, we created a stimulus display where circular dots (each subtending 0.03 degrees of visual angle) moved linearly in a random direction at a speed of 1.2 degrees of visual angle per second on a gray background. The direction of motion of each dot was idiosyncratic to that dot. At any given time, ten of these dots (5 black and 5 white) were present on the screen. When a dot exited the screen, it was immediately replaced with one of the same color starting at a random position in a random direction. The screen was divided in the middle by a vertical green line. This stimulus was displayed for 31 seconds. Halfway through the stimulus display, an “unexpected moving object” (UMO) - a black (same hue as the black dots) isosceles triangle that covered the same area as an individual moving dot - entered the screen on the right edge and traversed the entire screen in the center (bisecting the vertical green line in the middle) before exiting on the left side (Figure 2a).

The UMO moved at a speed that was randomly chosen from the following 12 speed conditions: 1, 1.75, 2, 2.42, 3, 3.5, 4, 4.58, 5.25, 6, 6.88, 8x. These speeds were relative to the speed of the dots, so the UMO moved minimally at 1.2 deg/s and maximally at 9.6 deg/s.

This stimulus display was created using MGL (Gardner Laboratory, Neuroscience Institute, Stanford University, 2014) with MATLAB (MATLAB R2013a, The Mathworks Inc., 2013). MGL is a MATLAB library designed for displaying full screen psychophysical stimuli. As MGL is not executable online, we created a library of functions - JGL - that consists of a JavaScript version of the corresponding MGL functions, allowing for online deployment of MGL code.

Participants were instructed to count the number of times the dots crossed the green vertical line. Each participant was assigned to one of two color conditions (counting crossings by either white or black dots only) and one of the 12 UMO speed conditions detailed above, yielding 24 (12x2) experimental conditions. Each participant performed only a single trial. The logic was the same as

in experiment 1: Once someone saw the unexpected stimulus, it is no longer an unexpected stimulus, so the nature of the task would be transformed to a visual search task in subsequent trials. Before the actual task, participants were exposed to a 20 second practice run where they were also asked to count the number of crossings, but no UMO entered the screen. The primary purpose of this practice run was to familiarize participants with the task.

Measures and Analysis

After completion of the task - in addition to reporting the number of observed crossings we asked participants to indicate whether they observed anything unexpected and if so, what. We also asked participants all the other questions we asked in experiment 1. The data were recorded in a sql-lite database via Psiturk (McDonnell et al., 2012). As in experiment 1, qualitative responses were coded by raters that were blind to the experimental condition as to whether the response indicated the detection of an unexpected object. All data were analyzed using MATLAB (Mathworks, Inc.), performing Chi square tests, due to the categorical nature of the dependent variable.

Results

There was an effect of endogenous attention on detection performance of unexpected objects. We conceptually replicated Simons & Chabris (1999) - participants were much more likely to detect the UMO when counting line crossings of black dots than when counting line crossings of white dots: $\chi^2(1, N = 1177) = 214.02, p < 0.001$; Cramer's $V = 0.43$.

There was an effect of speed on the probability that the unexpected object was detected, with faster objects being more likely to be detected (Figure 2). In contrast to experiment 1, this was the case for both attention conditions - whether participants were counting line crossings of white dots or black dots: white $\chi^2(11, N = 615) = 84.85, p < 0.001$; Cramer's $V = 0.37$; black $\chi^2(11, N = 562) = 23.11, p = 0.017$; Cramer's $V = 0.20$. However, the effect was much stronger for the condition in which participants counted line crossings of white dots.

-- Figure 2 about here --

Figure 2: a) Schematic of task and stimulus display. b) Proportion of participants reporting a relevant unexpected object, by experimental condition. x-Axis: Speed of object relative to counted dots. y-Axis: Proportion reporting seeing the unexpected object. Solid line: Participants reporting dot crossings by black dots. Dashed line: Participants reporting dot crossings by white dots. Error bars indicate standard error of the mean.

Discussion

As in experiment 1, we were able to show that the detection rate of the unattended UMO was dramatically modulated by speed such that it was more detectable for faster speeds.

In contrast to experiment 1, we found that only 20% of participants detected the slowest unattended (white) UMO - the one that moved with the same speed as the dots. The corresponding performance in experiment 1 for the slow unattended UMO was 60%. This result was in line with Simons & Chabris (1999), suggesting that participants in 2017 do expect the gorilla.

Also in contrast to experiment 1, there was an - ever so subtle - effect of speed on the detection of attended (black) UMOs as well. However, this could be accounted for by the fact that not all participants paid perfect attention to the task. Presumably, faster speeds attract more attention in general. In that case, a faster speed would be more detectable regardless of what people paid attention to.

One limitation of experiment 2 is that we did not know where the threshold of attentional override was, so we sampled a range of speed conditions that was at once too wide and too finely binned. Consequently, we are unable to tell whether the effect of speed on attentional capture is linear or whether there is a sharp threshold. Another fundamental limitation consists in the fact that we cannot tell whether this is an effect of absolute or relative speed, given these experiments and whether this effect is limited to black UMOs.

Experiment 3

The first two experiments confounded relative speed with absolute speed. In other words, faster moving UMOs were less similar to the attended objects. Of course, the critical role of similarity in the detection is well known (Most et al., 2001), with stimuli that contrast with the background being more salient and detectable than others, regardless of what that difference is, i.e., motion. To establish the speed dependence of UMO detection, we have to rule out object/background similarity as a parsimonious explanation of our data. To do so, we performed another experiment with UMOs that were moving slower than the background dots. In addition, we took this opportunity to address the question as to whether the speed dependence is particular to black UMOs or whether it generalizes to other colors.

Method

Participants

Participants were recruited through the Amazon “Mechanical Turk” platform. We recorded data from 2800 participants. We excluded participants using similar criteria as in experiments 1 and 2, retaining data from 1591 participants for further analysis. Specifically, we used the practice count as a filter, excluding data from participants with a count that was more than 25% off. This cutoff is arbitrary, but as in experiment 1 and 2, the results supported the same conclusions for a very broad range of such cutoffs. In general, we aimed for a tradeoff between retaining as much data as possible and excluding people who were unlikely to have been engaged in the task. We attribute this large attrition rate, given a relatively modest criterion for accuracy, to the challenging nature of the task in combination with the recently deteriorating quality of the mTurk participant pool (Kennedy et al., 2018). Participants were randomly assigned to one of 24 experimental conditions. As in experiments 1 and 2, all procedures were approved by the University Committee on Activities Involving Human Subjects (UCAIHS) at NYU.

Stimulus and Task

We used the same stimulus display we created for experiment 2, but with some modifications. Importantly, we extended the color conditions to both black and white UMOs, fully crossed with whether to count black or white dots. Participants were assigned to one of these 4 conditions, where they would be asked to count the dots of a target color (black or white) and could encounter one of two UMOs (black or white). In addition, we used a relative speed range that included slower as well as faster speeds. Specifically, there were 6 speed conditions: UMO moving at [$\frac{1}{4}$ $\frac{1}{2}$ 1 2 4 8] times relative to the speed of the dots.

To accommodate the lower speed conditions, allowing the UMO to fully traverse the screen in a reasonable time, we had to increase the speed of the dots to 1.8 degrees of visual angle per second and we extended the overall stimulus duration (how long the dots moved on the screen) to 37 seconds. The 6 speed conditions were fully crossed with the color attention conditions.

Measures and Analysis

All procedures are the same as in experiment 2, but in addition we separately analyzed the data from matching conditions (UMO color matched the color of the attended dots) and mismatching conditions (UMO color did not match the color of the attended dots). Also, we performed a bootstrap to estimate the statistical reliability of the asymmetry of speed slopes, resampling the absolute slope differences 1 billion times.

Results

The color of the UMO did not matter (Figure 3a). Detection rates for white UMOs and black UMOs were indistinguishable. What critically influenced detection rates was whether the attended color of the dots matched the color of the UMO. If the color of the unexpected object matched the attended color, detection rates were much higher (average detection rates: white/white: 58% +/- 5%, black/black: 67% +/- 5% and white/black as well as black/white both 23% +/- 4%, respectively). Thus, we collapsed these 4 conditions into two conditions – matching and mismatching conditions for further analysis.

-- Figure 3 about here --

Figure 3: a) The proportion of participants reporting a relevant unexpected object, by condition. x-Axis: Attention/UMO color conditions – W = white, B = black. The first letter in each combination represents the UMO-color and the second the color of the dots participants were instructed to count. y-Axis: Proportion of participants reporting an unexpected object. b) The proportion of participants reporting a relevant unexpected object, as a function of speed. x-Axis: Speed of UMO relative to speed of counted dots. y-Axis: Proportion of participants reporting the unexpected object. Solid line: Participants presented with an UMO matching the color of the dots they were instructed to count (i.e. black/black and white/white). Dashed line: Participants presented with an UMO of a different color than the dots they were instructed to count (white/black and black/white). Error bars indicate 95% confidence interval.

We replicated the effects of experiment 2: speed was a major determinant of detectability in mismatching conditions (Figure 3b). In addition, increased speed enhanced the detectability of UMOs in matching as well as mismatching conditions, although the profiles as to how speed impacted detectability differed between matching versus mismatching conditions: In the matching color condition, detection rates range from 44% for $\frac{1}{4}$ speed through 60% for matching speed to 78% for 8x speed. In the mismatching color condition, detection rates range from 12% for $\frac{1}{4}$ speed through 10% for matching speed to 51% for 8x speed. One interesting difference pertains to the effects of speed in the matching conditions. In experiment 2, performance in the matching condition was close to ceiling, where this task was much harder. Thus, we are able to show an effect of slower speeds on detection, if participants are expecting objects in that color, i.e. in matching conditions. In the mismatching conditions, only high speeds have an effect on detection. Plausibly, attention has to be captured first, if the unexpected object color does not match the feature people are attending to, which happens at faster speeds.

Slower speeds of mismatching UMOs were slightly more detectable than conditions in which the UMO speed was the same as that of the dots. However, the slope from $\frac{1}{4}$ to 1 (relative speed) was much shallower than that of from 1 to 4 (relative speed): bootstrap test $p = 1.62e-6$. In fact, this might underestimate the real difference in slope, as the slope from 2x relative speed to 8x relative speed was even steeper ($p = 0.0039$).

Discussion

We replicated the key effects from experiments 1 and 2, showing that speed dependence of UMO detection generalizes to UMO colors other than black. In addition, we ruled out a critical confound of experiment 2, namely the similarity of the UMO to the dots.

The effect of speed in the mismatched condition was quite dramatic – in the highest speed condition we tried, speed completely overrides selective attention, as in that condition, the majority of participants see the UMO, which is higher than the detection rate for matching conditions at slow speed.

Note that overall UMO detection levels in experiment 3 are somewhat lower than in experiment 2, which we attribute to the fact that the task was somewhat harder – we increased the speed of the dots to count by 50% and participants were asked to count for longer, which was reflected in participant comments mentioning task difficulty. In addition, it has been shown that the quality of the mTurk pool has degraded recently (Kennedy et al., 2018), although this is controversial (Johnson & Ryan, 2018).

UMO detection rates were lowest when the speed of the UMO matched that of the dots, as would be expected from a similarity-based account. However, simple probability summation (Verghese et al., 1999) could also account for this effect, and it is hard to tease apart a similarity from a probability summation interpretation on the basis of existing data. Simply put, in the slowest condition, the UMO is on the screen four times as long as in the condition where the UMO speed matches that of the background dots, so participants could be expected to be more likely to see it, just from random sampling of the screen. Regardless, we found strong support for the notion that faster relative speeds enhance detectability much more than slower relative speeds. It is remarkable that this is the case, given how briefly the fast UMOs are on the screen (the 8x speed UMO is on the screen only $\frac{1}{32}$ of the time that the $\frac{1}{4}$ x speed UMO is).

It would have been interesting to see whether a very fast UMO could fully close the detectability gap between matching and mismatching conditions. Alas, this is not practical, given current technology. Most of our participants will be running the experiment from LCD screens, with refresh rates in the 50-60Hz range. On the basis of our data, a stimulus could be expected to close that gap around a

relative speed of 16x. Given the base speed of about 1.8 deg/sec and given the spatial frequency composition of the stimuli, the temporal frequencies needed to achieve 16x relative speed would exceed Nyquist frequency (Nyquist, 1928), given typical refresh rates. Consequently, the stimuli might be temporally aliased, and we would not know if any given mTurk observer was seeing flicker or the actual stimulus. Thus, we recommend follow-up studies aimed at determining whether fast relative speeds can fully overcome the effects of attentional set to be performed in a laboratory setting, under controlled conditions using displays with sufficiently high refresh rates, around 120 Hz or higher (Poth et al., 2018). The primate motion detection system extends well into this speed range – well suited to detect stimuli moving at 32 deg/sec - and beyond, particularly in the periphery (Nover et al., 2005).

General Discussion and Conclusions

It has been suggested that unattended, unexpected objects are hard to see, even if they are moving (Simons & Chabris, 1999). This has become the dominant view in the field in the past 20 years (Galotti, 2013). We did three experiments - a direct replication and two conceptual extensions - to revisit this issue in light of the fundamental plight of organisms facing an environment of potential predatory uncertainty. Our results from these experiments consistently indicate that it is in fact the case that unexpected moving objects (UMOs) are less detectable if they are unattended. However, this effect strongly depends on the speed of the unattended UMO. If the UMO speed is slow, detection rates of UMOs in the unattended color are low, whereas they approach the detection rates of UMOs in the attended color if UMO speed is high relative to the speed of task-relevant objects.

We take these results to indicate the existence of a speed-based sentinel system that dynamically allocates attentional resources and overrides endogenous attention if an unexpected object is detected by this system. The high speed of the unexpected object could signify that it might pose an existential threat to the organism, i.e., the presence of a predator and therefore worth attending, regardless of the task that the organism was engaged in up to that point. These findings are in line with recent research suggesting that animacy is a critical aspect of the detection of unattended objects (Calvillo & Jackson, 2014; Gao & Jia, 2016), although our account is purely bottom-up, as it is driven solely by object speed.

This interpretation has a broader range of implications. If our interpretation of these results is correct, the commonly accepted view of change blindness effects might be misleading. It is not necessarily a problem for the organism that people do not readily perceive slowly moving unexpected objects (Simons & Chabris, 1999). It could be argued that the system deploying endogenous attention is working well (participants were doing the task as instructed) and the speed of the UMO was not high enough to justify an override by the sentinel system. If the sentinel system was triggered by any UMO, the false positive rate would presumably be so high as to be rendered useless, as we perhaps see in some attentional disorders.

Even more broadly, much of contemporary experimental psychology implies a fundamentally incompetent organism: People are not able to perceive things reliably, let alone veridically (Bach &

Poloschek, 2006), they are not able to pay attention (Simons, 2000b), they have unreliable memory (Loftus & Palmer, 1996), they make bad decisions (Kahneman, 2011) and they are unable to learn from them (Guilbault et al., 2004). Conversely - and we believe this view is supported by the data presented in this paper - one could argue that any cognitive system faces a challenging optimization problem as it is necessarily operating with limited resources in the face of profound existential uncertainty. Testing this system in laboratory settings might not constitute a fair test of what it is designed to do, so results gained in this way might be misleading.

But given that all organisms must perform survival relevant tasks, e.g., foraging, within a fundamentally uncertain environment, it is not implausible that the ones who did so successfully would have evolved mechanisms to manage this uncertainty while at the same time allowing them to focus on the task at hand. Given how much cortical real estate primate brains devote to motion processing (Wallisch & Movshon, 2008), it makes sense to utilize motion information in that way, allowing for the fast redeployment of attention.

We believe the results presented in this study strongly support this view, but there are several limitations of our research, most of which are due to technical limitations inherent in the study design in conjunction with the nature of recording data online. For instance, it would have been interesting to see whether the effect of speed eventually saturates or whether it is strong enough to completely match the effect of endogenous attention. One might also wonder how changes in motion direction, not speed, affect detection performance (Howard & Holcombe, 2010). Unfortunately - due to the nature of this task - we can only use a single trial per participant, so we were unable to fully explore this large parametric space in this study. These issues are not crucial to our theoretical point, but it would be interesting to know which of these possibilities is the case empirically.

To summarize, we conclude the likely existence of a speed-based sentinel system overriding task-dependent endogenous attention. This system would aid organisms in the management of resource allocation when pursuing survival relevant tasks in the face of fundamental uncertainty. Perhaps adopting such a hedging approach would be a reasonable strategy to help minds to cope with an uncertain world. Such a speed-based sentinel system to dynamically redeploy attention would constitute another example of how cognitive systems - implicitly mindful of their own limitations - match limited resources to task demands. This seems to be a strategy widely implemented in the brain where several competing demands have to be met by the system at once, e.g. the retina in the visual system. For example, organisms with a fovea effectively have two visual systems - one (central) with high resolution for shape and the other high sensitivity for motion (peripheral), which necessitates making rapid eye movements to dynamically sample the environment at high resolution. These results also profoundly challenge common wisdom about the nature of attention:

“It is a well-known phenomenon that we do not notice anything happening in our surroundings while being absorbed in the inspection of something” (Balint, 1907). Whereas this has been believed for 100 years, and emphasized for 20 (Chabris et al., 1999), we can show that this is demonstrably not the case, if the “anything” happens to be rapidly moving, relative to background.

Finally, it has not escaped our notice that this framework - implying a dynamic tug-of-war between endogenous and exogenous attention - potentially has direct applications for a deeper understanding of attention deficit disorders, which is still wanting. If it is the case that it is this sentinel system that is malfunctioning in people with attention disorders, there are different conceivable possibilities that all lend themselves to be tested empirically. For instance, someone who suffers from an attention deficit disorder might simply have a threshold of attentional capture that is too low, yielding too many false positives.

References

- Abrams, R. A., & Christ, S. E. (2003). Motion onset captures attention. *Psychological Science*, 14(5), 427-432.
- Bach, M., & Poloschek, C. M. (2006). Optical illusions. *Advances in Clinical Neuroscience and Rehabilitation*, 6(2), 20-21.
- Barlow, H. B. (1953). Summation and inhibition in the frog's retina. *The Journal of physiology*, 119(1), 69-88.
- Calvillo, D. P., & Jackson, R. E. (2014). Animacy, perceptual load, and inattention blindness. *Psychonomic Bulletin & Review*, 21(3), 670-675.
- Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the acoustical society of America*, 25(5), 975-979.
- Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *The Journal of physiology*, 187(3), 517-552.
- Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 353(1373), 1245-1255.
- Drew, T., Võ, M. L. H., & Wolfe, J. M. (2013). The Invisible Gorilla Strikes Again: Sustained Inattention Blindness in Expert Observers. *Psychological science*, 24(9), 1848-1853.
- Drews, F. A., Yazdani, H., Godfrey, C. N., Cooper, J. M., & Strayer, D. L. (2009). Text messaging during simulated driving. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 51(5), 762-770.
- Galotti, K. M. (2013). *Cognitive psychology in and out of the laboratory*. Sage Publications.
- Gao, H., & Jia, Z. (2016). Detection of Threats under Inattention Blindness and Perceptual Load. *Current Psychology*, 1-7.
- Goldstein, E. B., & Brockmole, J. (2016). *Sensation and perception*. Cengage Learning.
- Guilbault, R. L., Bryant, F. B., Brockway, J. H., & Posavac, E. J. (2004). A meta-analysis of research on hindsight bias. *Basic and Applied Social Psychology*, 26(2-3), 103-117.
- Hall, J. R., Cuthill, I. C., Baddeley, R., Shohet, A. J., & Scott-Samuel, N. E. (2013). Camouflage, detection and identification of moving targets. *Proceedings of the Royal Society of London B: Biological Sciences*, 280(1758), 20130064.
- Hillstrom, A. P., & Yantis, S. (1994). Visual motion and attentional capture. *Attention, Perception, & Psychophysics*, 55(4), 399-411.
- Holcombe, A. O. (2009). Seeing slow and seeing fast: two limits on perception. *Trends in cognitive sciences*, 13(5), 216-221.
- Howard, C. J., & Holcombe, A. O. (2010). Unexpected changes in direction of motion attract attention. *Attention, Perception, & Psychophysics*, 72(8), 2087-2095.
- James, W. (1890). *The principles of psychology*. Read Books Ltd.
- Jaschinski, W. (2002). The proximity-fixation-disparity curve and the preferred viewing distance at a visual display as an indicator of near vision fatigue. *Optometry and Vision Science*, 79(3), 158-169.

Johnson, D. and J. Ryan (2018). Amazon mechanical turk workers can provide consistent and economically meaningful data. Munich Personal RePEc Archive No. 88450.

Kahneman, D. (2011). Thinking, fast and slow. Macmillan.

Kawahara, J., Yanase, K., & Kitazaki, M. (2012). Attentional capture by the onset and offset of motion signals outside the spatial focus of attention. *Journal of vision*, 12(12), 10-10.

Kennedy, R., Clifford, S., Burleigh, T., Waggoner, P., & Jewell, R. (2018). The Shape of and Solutions to the MTurk Quality Crisis. Available at SSRN.

Kreitz, C., Furley, P., & Memmert, D. (2016). Inattention blindness is influenced by exposure time not motion speed. *The Quarterly Journal of Experimental Psychology*, 69(3), 495-505.

Loftus, E. F., & Palmer, J. C. (1996). Eyewitness testimony. In *Introducing Psychological Research* (pp. 305-309). Macmillan Education UK.

Rock, I., Linnett, C. M., Grant, P., & Mack, A. (1992). Perception without attention: Results of a new method. *Cognitive psychology*, 24(4), 502-534.

Maunsell, J. H. (2015). Neuronal mechanisms of visual attention. *Annual Review of Vision Science*, 1, 373-391.

McDonnell, J.V., Martin, J.B., Markant, D.B., Coenen, A., Rich, A.S., and Gureckis, T.M. (2012). psiTurk (Version 1.02) [Software]. New York, NY: New York University. Available from <https://github.com/NYUCCL/psiTurk>

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 inattention blindness. *Psychological science*, 12(1), 9-17.

Müller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, 285(5432), 1405-1408.

Nover, H., Anderson, C. H., & DeAngelis, G. C. (2005). A logarithmic, scale-invariant representation of speed in macaque middle temporal area accounts for speed discrimination performance. *Journal of Neuroscience*, 25(43), 10049-10060.

Nyquist, H. (1928). Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2), 617-644.

O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of 'mudsplashes'. *Nature*, 398(6722), 34-34.

Pashler, H. (1992). Attentional limitations in doing two tasks at the same time. *Current Directions in Psychological Science*, 1(2), 44-48.

Pashler, H. (1994). Dual-task interference in simple tasks: data and theory. *Psychological bulletin*, 116(2), 220.

Pelli, D. G. (2008). Crowding: A cortical constraint on object recognition. *Current opinion in neurobiology*, 18(4), 445-451.

Pelli, D. G., & Tillman, K. A. (2008). The uncrowded window of object recognition. *Nature neuroscience*, 11(10), 1129-1135.

Posner, M. I. (1980). Orienting of attention. *Quarterly journal of experimental psychology*, 32(1), 3-25.

- Poth, C. H., Foerster, R. M., Behler, C., Schwanecke, U., Schneider, W. X., & Botsch, M. (2018). Ultrahigh temporal resolution of visual presentation using gaming monitors and G-Sync. *Behavior research methods*, 50(1), 26-38.
- Priebe, N. J., Churchland, M. M., & Lisberger, S. G. (2002). Constraints on the source of short-term motion adaptation in macaque area MT. I. The role of input and intrinsic mechanisms. *Journal of neurophysiology*, 88(1), 354-369.
- Rempel, D., Willms, K., Anshel, J., Jaschinski, W., & Sheedy, J. (2007). The effects of visual display distance on eye accommodation, head posture, and vision and neck symptoms. *Human factors*, 49(5), 830-838.
- Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in cognitive sciences*, 1(7), 261-267.
- Simons, D. J. (2000a). Attentional capture and inattention blindness. *Trends in cognitive sciences*, 4(4), 147-155.
- Simons, D. J. (2000b). Current approaches to change blindness. *Visual cognition*, 7(1-3), 1-15.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28(9), 1059-1074.
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in cognitive sciences*, 9(1), 16-20.
- Strayer, D. L., Drews, F. A., & Crouch, D. J. (2006). A comparison of the cell phone driver and the drunk driver. *Human factors*, 48(2), 381-391.
- Taleb, N. N. (2007). *The black swan: The impact of the highly improbable* (Vol. 2). Random house.
- Treue S, Maunsell JH. 1996. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382: 539-41
- Thompson, B., Hansen, B. C., Hess, R. F., & Troje, N. F. (2007). Peripheral vision: Good for biological motion, bad for signal noise segregation?. *Journal of Vision*, 7(10), 12-12.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive psychology*, 12(1), 97-136.
- Verghese, P., Watamaniuk, S. N., McKee, S. P., & Grzywacz, N. M. (1999). Local motion detectors cannot account for the detectability of an extended trajectory in noise. *Vision research*, 39(1), 19-30.
- Wallisch, P., & Movshon, J. A. (2008). Structure and function come unglued in the visual cortex. *Neuron*, 60(2), 195-197.
- Wandell, B. A. (1995). *Foundations of vision*. Sinauer Associates.