THE BETTER TO SEE: THE ROLE OF SUPPRESSION IN INATTENTIONAL BLINDNESS

BY

KATHERINE WOOD

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Psychology
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2020

Urbana, Illinois

Doctoral Committee:

Professor Daniel J. Simons, Chair
Professor Diane M. Beck
Professor Aaron S. Benjamin
Professor John E. Hummel
Professor David E. Irwin

ABSTRACT

There is a well-documented disconnect between the amount of information we subjectively feel that we have access to and the amount of information our visual system can actually encode, evaluate, and maintain. The former feels immense and effortlessly retrieved; the latter, we know from decades of neural and behavioral work, is quite limited. We preserve a feeling of visual detail and richness in spite of our capacity-limited system by the process of selective attention. We enhance the representation of information to which we attend and inhibit the representation of unwanted information. In so doing, we mitigate the deluge of information coming into the system and dedicate our limited resources to processing only that which we have selected. One of the most remarkable impacts of this deployment of selective attention is the phenomenon of inattentional blindness, in which we fail to detect obvious stimuli or events because our attention is focused on some task. This dissertation describes a series of experiments that use inattentional blindness to explore the impact of suppressing unwanted stimuli and reveal the extensive reach of the inhibitory aspect of selective attention.

The series of experiments outlined in Chapter 2 investigated the scope of inhibition and selection by utilizing displays in which one group of objects was homogenous with respect to the task-critical feature, in this case color, while the other group of objects varied either within or across trials. When subjects ignored the heterogeneous objects, they exhibited near-total inattentional blindness for unexpected objects in actively ignored colors, previously ignored colors, and completely novel colors not present in the display at all. When subjects attended to these variable objects, only unexpected objects that matched the attended color on the critical trial attracted notice at any appreciable rate. Suppression acts broadly to inhibit irrelevant

information, even if it has not been encountered in the context of the current task. Selection, by contrast, is narrow and confined to the immediate demands of the task.

The experiments described in Chapter 3 suggest a similar tendency for broad suppression of space, in addition to features. Subjects played a game requiring them to shuttle objects back and forth across a display while avoiding hazards. The game environment contained a large amount of visual information, but only a small subset of the display was task-critical. Subjects tended to detect unexpected objects that occurred in the most task-relevant areas of the display, with attention appearing to be concentrated around their avatar in the area where the risk of collision with a hazard was highest. Noticing dropped off dramatically with relatively small distances from this locus of attention, even though the object in question was a novel color in the display. Subjects automatically prioritized the most critical areas of the display and suppressed information in the rest.

The results in Chapter 4 clarify the time course of when noticing occurs in inattentional blindness. Subjects showed only small increases in their likelihood of noticing an unexpected object even as the amount of time it spent in the display doubled or tripled, and they consistently identified the object as having appeared near its onset point when asked to report where it was when they first detected it. This suggests that unexpected objects have a relatively brief period of time after they onset to break through suppression, after which they are unlikely to draw notice at all.

Together, the results in these chapters reveal the power of suppression. It acts broadly and quickly to inhibit task-irrelevant information, and plays a dramatic role in shaping what information reaches awareness.

ACKNOWLEDGEMENTS

I owe profound thanks to my advisor Daniel Simons. Working with a scientist and thinker of his caliber has been a singular privilege, and has made me not just a better scientist and thinker myself, but a better speaker, writer, and teacher.

I would also like to thank my committee, John Hummel, Diane Beck, Dave Irwin, and Aaron Benjamin, not only for their insights and recommendations on my dissertation work, but throughout the years in brownbags, lab meetings, and informal conversations. My work been consistently improved by their comments and ideas.

I am also grateful to both the University of Illinois and the National Science Foundation for supporting my research with fellowships for the duration of my graduate studies. Guaranteed funding is a rare and liberating gift that I was immensely lucky to receive and for which I am deeply appreciative.

Thank you to everyone in the Attention and Perception division for advice and feedback. I am especially grateful for my fellow graduate students. A burden shared is not always a burden halved, but their support certainly made it easier to carry.

I am grateful to my family for their unfailing support, encouragement, and interest in my work. It has meant so much to me to know they are proud of me.

Finally, thank you to Jason Rock, for everything.

TABLE OF CONTENTS

CHAPTER 1: BACKGROUND

A CAPACITY-LIMITED SYSTEM

Our experience of the world is a grand illusion. Our perception feels expansive and encompassing, taking in all the details of the space around us with speed and ease. We need only move our eyes to bring any object or space we desire under the powerful beam of our attention. We behold it, and it is ours.

The perceptual system that enables this illusion is not as capacious as our subjective experience might lead us to believe. What feels limitless can in fact be profoundly limited, and what feels rich and detailed can be quite coarse. In actuality the perceptual system is constrained at every level, beginning early in visual cortex. Stimuli compete for representation as early as area V4 in visual cortex, and having multiple stimuli in the receptive field of a neuron will suppress its response to its preferred stimulus (Desimone & Duncan, 1995). This appears to occur in later areas of visual cortex as well (Beck & Kastner, 2009). Deploying attention to one of the competing stimuli will resolve the competition in favor of the attended object, but only one; competition cannot be biased in favor of multiple items as effectively as one (Scalf & Beck, 2010).

Severe limits on capacity, processing, and memory similarly pervade the visual system at later levels. While we seem to have access to a tremendous amount of visual detail, this persists for only a few hundred milliseconds before rapidly decaying (Sperling, 1960). Once this level of detail plateaus, we appear to be able to hold at most about four visual objects in working memory at one time (Cowan, 2001; Rouder, Morey, Cowan, Zwilling, Morey, & Pratte, 2008), although this number can be even lower for more complex objects (Alvarez & Cavanagh, 2005).

Our visual store is limited not only in the absolute amount of information it can actively maintain, but also in the speed with which it can operate. In the phenomenon known as the *attentional blink*, subjects are asked to monitor a stream of rapidly presented stimuli for some set of targets; if two targets appear within 500 ms of each other, detection of the second target suffers dramatically as a result of attending to the first target (Raymond, Shapiro, & Arnell, 1992). While there is some debate about whether the attentional blink reflects a limit on encoding items into visual working memory or the speed of attentional disengagement and redeployment (see Marois & Ivanoff, 2005), there is no doubt that it demonstrates a profound limit on the rate of information processing in the system.

The phenomenon of *change blindness* delivers an even more striking demonstration of limits in the system. Instead of impaired detection of a single, briefly-presented target, change-blind subjects miss large changes to their visual environment when these changes coincide with a visual interruption or transient (Grimes, 1996; Rensink, O'Regan, & Clark, 1997). Changes do not have to be subtle to go unnoticed. Subjects fail to notice when every element in a display changes color, so long as the global statistics are preserved (Saiki & Holcombe, 2012). This change blindness does not seem to arise merely from an impoverished representation of the visual environment, as subjects demonstrate memory for the objects in the scenes they are monitoring (Hollingworth & Henderson, 2002) and in some cases even have an explicit representation of the identity of an object before it changed (Simons & Levin, 1998; Mitroff, Simons, & Levin, 2004). Thus change blindness appears to demonstrate a limit of information integration and comparison (Blackmore, Brelstaff, Nelson, & Trościanko, 1995; Hollingworth, 2003).

SELECTIVE ATTENTION

We observe capacity limits at every level of the visual system, from the quality of the representation in visual cortex to the encoding speed, storage capacity, and integration of visual information. With these limits in place, how are we able to parse the glut of information out in the environment and make sense of our visual world? Selective attention is thought to be the primary means by which we carve understanding from the flood of visual information, and its effects extend from the earliest levels of visual cortex to the highest levels of visual processing (Desimone & Duncan, 1995; Carrasco, 2011).

The effects of attention are apparent beginning at the level of the neuron. The responses of neurons to their preferred stimuli are enhanced when subjects attend to that stimulus. This can be observed as early as V1 (McAdams & Maunsell, 1999), V4 (Spitzer, Desimone, & Moran, 1988), and up through MT in monkeys (Treue & Trujillo, 1999). Selective attention also resolves the competition in favor of the selected item when multiple items are competing for representation, observable at the level of the firing activity of the neuron in monkeys (Desimone & Duncan, 1995) and the BOLD response in humans (Beck & Kastner, 2009; Scalf & Beck, 2010).

This enhancement can be observed beyond neural responses to stimuli. Selectively attending to a stimulus can enhance its apparent contrast (Carrasco, Ling, & Read, 2004) and brightness (Tse, 2005). Directing our attention to a particular location in space speeds responses to stimuli that appear there (e.g. Posner, 1980). We can similarly selectively attend to a particular feature, such as motion direction (Saenz, Buracas, & Boynton, 2002) or color (e.g. Brawn &

Snowden, 1999), although feature-based attention is not as consistent in its effects and is more sensitive to the particulars of the task than spatial attention (see Liu, Stevens, & Carrasco, 2007).

At a higher level, attention allows us to perform demanding tasks. We can track multiple dynamically-moving objects, even when they are identical to non-targets (Pylyshyn & Storm, 1988; Scholl, 2009) or continuously monitor several changing objects at once (Alvarez & Cavanagh, 2005). Outside of the laboratory, we execute complex and demanding attentional tasks as a matter of course. We find a friend in a crowd of a hundred faces, we drive, we navigate unfamiliar areas, we play fast-paced video games. Selective attention allows us to sift through the noise and allocate our limited resources effectively to accomplish these feats.

INHIBITION OF IRRELEVANT INFORMATION

If selection allows us to enhance the desired information in a busy environment, what happens to all of the information we do *not* wish to select? If competition for representation and therefore limited processing resources is inevitable, how do we ensure that the noise and bright confusion of a busy world does not overwhelm us? It seems that we accomplish this by deliberately dimming the house lights so that the spotlight of selection can shine out in the dark (Posner, Snyder, & Davidson, 1980).

Just as there are brain circuits devoted to selective attention, and the effects begin early in visual cortex, so too are there systems for suppressing unwanted information. Suppression of incoming information occurs extremely early, even prior to cortex; the LGN modulates responses to ignored stimuli, dampening the signal before it even reaches visual cortex (O’Connor, Fukui, Pinsk, & Kastner, 2002), and the basal ganglia and thalamus have been implicated in even earlier sensory gating (Nakajima, Schmitt, & Halassa, 2019). Ignoring of unwanted information is not controlled exclusively at the level of cortex, but like attentional selection is controlled in a variety of brain regions beginning in early sensory areas. Inhibition of distractor features can also act rapidly in visual processing, even before enhancement of target features (Moher, Lakshmanan, Egeth, & Ewan, 2014).

The suppression of ignored stimuli has profound perceptual and behavioral effects. In early studies of dichotic listening, in which subjects listened to one of two simultaneous audio streams, subjects routinely failed to encode content in the ignored channel (Moray, 1959; Treisman, 1964a) and failed to detect changes to the ignored channel such as a change in the language spoken (Treisman, 1964b) or a sudden shift to backwards speech (Cherry, 1953; Wood & Cowan, 1995). In analogous studies using superimposed visual displays, subjects were similarly hazy on the details of the ignored display and failed to notice odd changes to the action of the ignored video (Neisser & Becklen, 1975). This effect emerges for static pictures as well; people who attended to one of two superimposed pictures and ignored the other had no detailed representation of the ignored shape other than its most basic features, even when that shape was familiar (Rock & Gutman, 1981). These effects can also be observed at the level of the neural response. When people ignored strings of letters at fixation and attended to superimposed pictures, the BOLD response was the same regardless of whether people were ignoring meaningful words or meaningless strings of letters (Rees, Russel, Frith, & Driver, 1999).

In these cases, subjects were at least nominally aware that some kind of stimulus was present, even if they could not describe many of its features. Filtering of unwanted information can have far more extreme consequences than degraded or coarse representations; subjects can be completely unaware that any stimulus was ever present at all. In studies of *inattentional blindness*, people are completely unaware of novel and highly visible objects or events.

THE CONTRIBUTIONS OF INATTENTIONAL BLINDNESS

The same selective looking paradigm used to demonstrate unawareness of ignored visual content also showed that subjects failed to notice a woman with an umbrella walking through the middle of the ball game they were monitoring (Neisser & Dube, 1978; Becklen & Cervone, 1983). The failure to notice the umbrella woman demonstrated a phenomenon distinct from the sparse representations of ignored stimuli. Subjects were not actively ignoring her at all, yet she failed to draw notice as she appeared and moved through the middle of the attended action. This remarkable obliviousness is not induced exclusively by busy, dynamic displays teeming with distracting stimuli, either. Subjects routinely failed to notice a black square appearing in an otherwise empty display when they were attending to a central cross (Mack & Rock, 1998). This demonstration was particularly striking because displays did not contain distractors that needed to be actively filtered, but it nevertheless seemed that the mere act of attending to the cross made subjects miss the appearance of a new object.

Inattentional blindness provides a unique window into the functioning of selective attention and distractor inhibition because of the nature of the unexpected stimulus. With stimuli that are present in the display but to be ignored, subjects have an opportunity to direct attention to them, even if only to evaluate and filter them. While this does not invalidate the conclusions that have previously been drawn about the status of ignored items, it does mean that one cannot definitively say that the results arise from an absence of attention (Mack, Tang, Tuma, Kahn, & Rock, 1992). Inattentional blindness, by contrast, does allow us to gauge a response to a truly unattended stimulus. Subjects do not expect a novel object to appear, and have no opportunity to allocate attention to it or process it in any way until it appears in the display. Unexpected objects are therefore “uncontaminated” by any ongoing attentional processes, and can measure the attentional state of the system without influencing it. This property provides another way to investigate attentional processes, and an unobtrusive way to measure attention.

While many of the results of inattentional blindness experiments echo previous findings in other paradigms, some results have been unexpected. Many inattentional blindness experiments have demonstrated the dissociation between location of gaze and location of covert attention by repeatedly finding that subjects fail to notice unexpected objects even when they stare directly at them (Mack & Rock, 1998; Koivisto, Hyönä, & Revonsuo, 2004; Beanland & Pammer, 2010). Subjects do notice unexpected objects more often when they are close to the location of attention, however (Mack & Rock, 1998; Most, Simons, Scholl, & Chabris, 2000; Stothart, Boot, & Simons, 2015). These results are consistent with the body of work exploring covert attention (for a review see Carrasco, 2005), but go even further in showing that if we are not prepared for an object and it does not appear in the locus of attention, it can go unnoticed *entirely* and not merely suffer a reaction time or accuracy decrement.

While the results of inattentional blindness studies have reinforced the consistent results from studies of spatial attention, they have also clarified effects not observed as consistently. One such effect is that of feature-based attention; some studies find no benefit in terms of accuracy or reaction time for attending based on features, using paradigms such as visual search (Moore & Egeth, 1998; Shih & Sperling, 1996) and cueing (Theeuwes, 1989). While not all studies of feature-based attention reach this conclusion (e.g. Liu, Stevens, & Carrasco, 2007), the effect seems more elusive than the consistent effects of spatial attention. Inattentional blindness, in contrast to other paradigms, consistently finds robust effects of feature-based attention.

A typical inattentional blindness experiment on feature-based attention has subjects perform a multiple-object-tracking task with two sets of objects, each with a different feature. Subjects attend to one set of objects, such as white Ts and Ls, and ignore the other set, such as black Ts and Ls, and an unexpected object with some feature passes through the display. Unexpected objects that share the critical feature with the attended set are routinely noticed, while objects that share a feature with the ignored set are almost never noticed. This effect appears with luminance (Most, Simons, Scholl, Jiminez, Clifford, & Chabris, 2001), shape (Most, Scholl, Clifford, & Simons, 2005), color (Goldstein & Beck, 2016; Drew & Stothart, 2016), and even semantic category (Most, 2013). While feature-based attention may not hasten visual search or show the same cuing advantages that spatial attention does, inattentional blindness reveals that it can make the difference between a new object reaching awareness or not.

Similarly to feature-based attention, inattentional blindness has helped recontextualize the sometimes opposed results of experiments on attention capture. There has been extensive debate about the precise nature of attention capture, what aspects of a stimulus can capture attention, and under what circumstances capture occurs (Folk, Remington, & Johnston, 1992; Bacon & Egeth, 1994; Theeuwes, 2004). Depending on the subjects' attention set, there are circumstances under which sudden onsets or colored objects will capture attention, operationalized in these studies as a slowdown in visual search times for a target when the capturing stimulus is present. Inattentional blindness adds an important dimension to this discussion. A subject's attention set certainly has tremendous influence over whether an object reaches awareness, consistent with findings from the attention capture literature. However, neither abrupt onsets nor highly salient stimuli appear to gain preferential access to awareness, as subjects are routinely inattentionally blind to these objects (Mack & Rock, 1998; Most, Simons, Scholl, Jiminez, Clifford, & Chabris, 2001; Most, Scholl, Clifford, & Simons, 2005). Thus the "capture" observed in these search tasks may not extend to more complex tasks or stimuli, or tasks other than visual search for a target; there may also be a dissociation between attention capture as typically operationalized and explicit awareness of a stimulus.

Inattentional blindness can offer more than just replicating and extending findings in the extant literature. It is also capable of addressing questions few other paradigms can answer. Virtually no other paradigm can claim to show the results of processing under true conditions of inattention. If subjects know something is present but are ignoring it, or if the event is extremely low probability but not entirely unexpected, then one cannot ensure that a stimulus or event was truly unattended or received no processing. Inattentional blindness can make critical and unique contributions to studies of preattentive processing or of processing that does not require attention (e.g. Moore & Egeth, 1997). Another unique feature of the paradigm is that it provides a means of measuring attention completely independently from the primary task, without using techniques such as MRI or EEG. This is vital for studies investigating behavioral consequences of attention in the context of a particular environment or scenario. Requiring an additional task to provide a measure of a manipulation, such as response time in a vigilance task while driving in a simulator, may alter the way subjects allocate their attention compared to performing the driving task on its own. With inattentional blindness, however, subjects need only perform and dedicate attention to the primary task.

Inattentional blindness is also remarkably flexible and can be induced with virtually any type of display or task, from sparse displays with just one object (Mack & Rock, 1998), to moving groups of simple shapes (Most, Simons, Scholl, Jiminez, Clifford, & Chabris, 2001), to videos of real action (Neisser & Becklen, 1975; Simons & Chabris, 1999), to real-world scenarios (Hyman, Boss, Wise, McKenzie, & Caggiano, 2010; Chabris, Weinberger, Fontaine, & Simons, 2011; Simons & Schlosser, 2017). Because this phenomenon is so robust to choice of stimuli, display, and the nature of the primary task, it can be easily adapted for a task or context of interest. There is little concern that a change to the timing, global contrast, or number of objects onscreen will attenuate or abolish the effect.

THE PRESENT EXPERIMENTS

Although inattentional blindness is a paradigm with limitless potential, there are many fundamental questions that remained unanswered. One such question concerns the nature of suppression, and the basis on which unexpected objects are blocked from reaching awareness. Typically, experiments that require subjects to ignore items only feature a single homogenous set to ignore. In the real world, however, the amount of information we must filter is diverse and vastly larger that the amount of information we select at any one time. How does the system handle this asymmetry? Chapter 2 addresses this question in a series of experiments that use a variant of the typical multiple object tracking task (Most et al., 2001). Rather than relying on two homogenous set of objects, one set of objects is heterogeneous with respect to the critical feature. When subjects have to ignore not one specific color, but several, what are the consequences for the to-be-ignored colors? What are the consequences for completely novel colors not present in the display? The answers to these questions will indicate how broadly suppression operates, particularly when we must contend with highly variable unwanted information.

Chapter 3 addresses a similar question, but for space instead of features. Do we automatically parse a display and narrow our attention only to the most relevant areas depending on the task, without requiring any explicit direction? What aspects of a task serve to shape our attentional priorities? In Chapter 3's experiments, subjects play a game that requires them to navigate through a busy display and avoid hazards. Only a small portion of the display is actually immediately task-relevant at any given time due to the structure of the task and the nature of movement subjects are allowed. Will subjects respond to this structure and allocate their attention accordingly, concentrating in the area of highest risk and ignoring the rest of the display? The experiments of Chapter 3 will provide an indication of how task demands shape the allocation of attention in the absence of any explicit direction, and whether we automatically establish an attention set for task-relevant areas based on the nature of the task alone.

Chapter 4 describes a series of experiments exploring the question of when unexpected objects reach awareness. Subjects perform a multiple object tracking task and are presented with an unexpected object that varies in how long it is visible in the display; as part of the post-critical-trial questionnaire, they also report where the unexpected object was in the display when they first noticed it. The pattern of noticing rates, combined with the location reports, will sketch the time course of how unexpected objects reach awareness. Is noticing triggered by a particular event, such as onset or offset? Does the unexpected object break through at random times? Does the likelihood of noticing increase with time? Understanding this time course will provide insight into the process of suppression in inattentional blindness, particularly into whether it is a relatively uniform process or if its consistency fluctuates with time.

These chapters explore three dimensions of suppression in inattentional blindness: features, space, and time. Together they sketch out a more complete picture of the role of suppression in inattentional blindness and perception more broadly.

CHAPTER 2: BROAD SUPPRESSION, NARROW SELECTION

This chapter describes previously published experiments.[1]

People can allocate attentional priority to a particular feature, such as the color green, and process anything sharing that feature more efficiently. This is sometimes referred to as directing *feature-based attention* to a feature (e.g. Störmer & Alvarez, 2014) or an *attention set* for a feature (e.g. Adamo, Pun, Pratt, & Ferber, 2008). These attention sets have been shown to enhance processing for a feature across the visual field, allowing selection of a feature even if its location is not known (Bichot, Ross, & Desimone, 2005).

The effects of attention sets have been demonstrated most dramatically in inattentional blindness paradigms. Generally speaking, unexpected objects which are not in a subject's attention set rarely reach awareness, while unexpected objects which are in the attention set routinely draw notice (Most et al., 2001; Most et al., 2005). This effect is restricted to spatially relevant areas; people do not appear to maintain attention sets for ignored portions of the display (Stothart, Simons, Boot, & Wright, 2019).

Why do unexpected stimuli that match the ignored objects fail to reach awareness? By the "filtering" account (e.g. Broadbent, 1958), ignored objects fail to reach awareness because they are actively filtered or suppressed early in selection on the basis of their features. Unexpected objects that share features with the ignored objects are therefore screened out by the filter subjects establish to inhibit task-irrelevant information. By the "information pickup" or "perceptual cycle" account (e.g. Neisser, 1976; Neisser, 1979), these objects are not actively

---

[1] Wood, K. & Simons, D. J. (2017). Selective attention in inattentional blindness: Selection is specific but suppression is not. *Collabra: Psychology, 3*(1). Published and reproduced under a CC-BY license (http://creativecommons.org/licenses/by/4.0/).

filtered or suppressed. Rather, they are simply “not picked up.” They do not need to be filtered out because no information about them enters the system in the first place, and attentional selection occurs only by means of actively picking the desired information out of the environment (Neisser & Becklen, 1975).

To distinguish between these two broad accounts, one would have to show a difference between information that is being actively suppressed and information that is not being selected. Such a difference would indicate that some information is being actively filtered, while a failure to demonstrate such a difference would indicate that there is no difference between supposedly inhibited information and information that is simply not actively selected. Inattentional blindness is a particularly useful paradigm to address this question, because it allows for a distinction between actively ignored items and truly unattended items. In the following experiments, this property is leveraged to demonstrate that the same object, with the same features and degree of task-irrelevance, gains access to or is blocked from awareness based on whether it is actively inhibited or merely unattended.

INTRODUCTION

Inattentional blindness is a remarkable illustration of our ability to focus attention exclusively on relevant information and filter distracting information from awareness. Sustained inattentional blindness tasks induce these attentional demands by tasking subjects with tracking one set of moving objects and ignoring the rest. Subjects in this state are often oblivious to unexpected stimuli, even strange ones that remain in view for seconds at a time (Neisser & Becklen, 1975; Simons & Chabris, 1999).

Our likelihood of noticing an unexpected object is heavily influenced by its similarity to other objects in the display. If we attend to objects based on color, we typically notice objects that match the attended color and miss those that match the ignored color (Most, Simons, Scholl, Jiminez, Clifford, & Chabris, 2001). These similarity effects occur for other features, such as shape, and they are flexible: When subjects attended to black circles and squares and ignored white ones, they noticed a black circle or square much more often than a white circle or square. When they viewed the same display but attended on the basis of shape instead (black and white circles versus black and white squares), circles of either luminance were noticed more often than squares (Most, Scholl, Clifford, & Simons, 2005). What seems to matter most to noticing is the unexpected object's similarity along the critical dimension that defines the attention set. Differences in irrelevant feature dimensions have less impact on noticing.

Such similarity effects have been observed for many feature dimensions, and they are thought to result from the attention sets subjects form during the task (Most et al., 2005). There is little agreement, however, in how these attention sets are structured or what aspects of similarity drive noticing. These attention sets might be entirely feature-based: We enhance features matching the attended set, drawing in objects that share those features and suppressing the distractor features. Consistent with this idea, when subjects attend to objects of one color and ignore two other sets of colorful objects—one nearer to the attended set in color space, and one much farther away—unexpected objects with colors similar to the attended set in color space are more likely to be noticed and those with colors close to the ignored ones are missed (Drew & Stothart, 2016).

Alternatively, attention sets might operate not on the specific features of the attended and ignored items, but in a more categorical way. Rather than enhancing “white” and suppressing “black,” the attention set might enhance “lighter” and suppresses “darker.” When display objects are red-orange and yellow-orange, and subjects are instructed to attend to the “redder” set, they unsurprisingly notice unexpected red-orange objects (Goldstein & Beck, 2016). However, they also notice extreme examples of the relation (red when attending to “redder” or yellow when attending to “yellower”) just as often as exact color matches. Although red is a better fit for the “redder” category, it also deviates more from the ignored, yellow-orange objects. Consequently, it is unclear whether the high rates of noticing result from greater similarity to the attended category or greater dissimilarity from the ignored one.

Attention sets might even operate at a semantic level. People are more likely to notice an unexpected block-face ‘E’ than a block-face ‘3’ when attending to letters and ignoring numbers, but they notice the 3 more often when attending to numbers and ignoring letters (Most, 2013). These two objects are nearly identical in their low-level features, but noticing rates differed based on whether they matched the attended and ignored semantic categories.

Each of these selection mechanisms may contribute to inattentional blindness, but few extant studies distinguish among them. Almost all inattentional blindness displays use two sets of objects that are homogenous with respect to the critical feature, dimension, or category. If objects are differentiated on color, subjects typically only have to attend white and ignore black. Yet, black and white objects will end up in the same groups if they are separated on absolute color value, luminance, relationships like “lighter” and “darker,” or broad categories like “light things” and “dark things.” Similarly, unexpected objects that vary along the same critical feature dimension will fall into one category or the other regardless of what truly determines “similarity”—dark gray is close to black relationally, in terms of its RGB value, its luminance, and its broader category, for example—making it difficult to determine what is driving the similarity relationship.

This redundancy can also also undermine conclusions about similarity—even when subjects must ignore two separate colors while attending one (e.g., Drew & Stothart, 2016), they could segment the objects not based on the actual color value, but according to “bright colors” and “dark colors,” or “hot” and “cold” colors. Red objects and red-orange objects are both “redder” than yellow-orange objects, but also featurally similar to each other and dissimilar from the ignored objects (Goldstein & Beck, 2016). To determine which aspect of similarity drives the effects of attention sets on noticing, an experiment must isolate each possible mechanism, eliminating other ways to parse the display.

To remove this redundancy, we employed stimuli that are not easily separable along a single dimension but that could be distinguished either by individuating all of the features of the objects or by grouping them into coarse categories. This approach allows us to separate the contributions of feature similarity and category similarity on noticing rates for unexpected objects.

EXPERIMENT 1

In order to examine how attention sets are formed, we used two sets of objects in Experiment 1. Subjects performed a multiple object tracking task in which they counted bounces for one set of objects in the display while disregarding the bounces of the other set. One set consisted of four white shapes and the other consisted four colorful shapes (black, red, yellow, and purple, with each color used once). These were chosen to prevent use of a single feature distinction to segment attended from ignored items (e.g., “white and black” or “light and dark”). To separate the objects, either subjects must track each object’s unique color (the “feature-based” hypothesis) or the objects must be sorted into coarser groups, such as “white” and “nonwhite” (the “category-based” hypothesis).

Experiment 1

	Unexpected Object Category			
	White	**Color**	**New**	**Total**
Attend White	90	75	98	263
Attend Nonwhite	108	99	83	290
Total	198	174	181	**553**

Experiment 2

	Unexpected Object Category			
	Critical Trial Color	**Previous Trial Color**	**New Color**	**Total**
Attend White	92	92	91	275
Attend Nonwhite	87	96	77	260
Total	179	188	168	**535**

Experiment 3

	Unexpected Object Category			
	Critical Trial Color	**Previous Trial Color**	**New Color**	**Total**
Attend Constant	65	56	59	180
Attend Variable	66	73	54	193
Total	131	129	113	**373**

Table 2.1. Number of subjects assigned to each attention and unexpected object condition after exclusions.

To determine which selection method people use, we can examine noticing rates for a unexpected green object that shares the category (“color” or “nonwhite”) but not the specific features (black, red, yellow, or purple) of that set. When people attend to the colorful shapes and ignore white, an unexpected green object should be noticed at high rates regardless of whether selection operates on categories of features. If selection is feature-based, a green object should be noticed because it is unique and salient. If selection is category-based, a green object should be noticed because it matches the category (“nonwhite") of the attended set. In both cases, green differs from the ignored shapes, and suppression of ignored shapes contributes to the likelihood of noticing (Most et al., 2001).

The informative case is when subjects are ignoring the colorful shapes. If selection is feature-based, then people should be ignoring “red, purple, black, and yellow.” Green differs from these colors, so it should escape suppression and be noticed at high rates. Alternatively, if selection is category-based and people are ignoring “nonwhite” shapes, then green falls within the suppressed category and it too should be suppressed (see Figure 2.1, column B).

Methods

Methods, procedures, target sample size, exclusion rules, stimuli, experimental code, and analysis scripts were preregistered sequentially, with each experiment preregistered before we started data collection for that experiment (https://osf.io/7pz35/). Data were analyzed using R (R Core Team, 2015) and are available on OSF. We report all data exclusions, measures, and manipulations here and in the preregistration (Simmons, Nelson, & Simonsohn, 2011).

Subjects. We aimed to collect usable data from 100 subjects in each of 6 conditions after exclusions (total target n = 600). Subjects were workers on Amazon Mechanical Turk who had at least 95% approval rates for their previously submitted HITs. We checked worker IDs against a database of prior subjects using TurkGate (Gideon & Goldin, 2013), and anyone who had previously participated in an inattentional blindness experiment from our laboratory or the laboratory of our collaborators was informed that they were not eligible for this HIT and were excluded prior to participating.

The need for signed consent was waived by the Institutional Review Board at the University of Illinois due to the low-risk nature of the experiment. Prior to the experiment, subjects were shown an information screen that provided experimenter and IRB contact information. It explained that their responses would be anonymous, described how their data would be used, and noted that their participation was voluntary.

Prior to analysis, we excluded subjects who skipped any questions during the experiment, miscounted in the tracking task by more than 50% in either direction on more than one trial, reported being younger than 18, reported experiencing any issues with the display or playback of the experiment, reported needing vision correction but not wearing it during the experiment, reported any form of colorblindness, or misidentified the number in Ishihara Plate 9 (Ishihara, 1990). Based on prior studies using similar exclusion rules and sampling from the same online population, we anticipated the need to exclude 40% of the subjects who completed the task.

Subjects were automatically recruited in batches of up to nine, with random assignment to the six conditions, until at least 1000 had completed the experiment, at which point no further batches were posted. In total, we recruited 1001 subjects. Subjects received $0.10 upon completing the experiment.

Materials and procedure. A demonstration of the experiment, identical to the one subjects completed (but without any data collection), can be found at http://simonslab.com/mot/set_demo.html.

The experiment was coded in Javascript, and was modeled on prior online sustained inattentional blindness tasks (Cary Stothart, personal communication, October 9, 2015; e.g. Drew & Stothart, 2016). At the start of the experiment, instruction screens informed subjects that they would see two sets of objects—a group of white shapes, and a group of nonwhite shapes—bounce around inside a blue rectangle. They were told to count how many times either the white

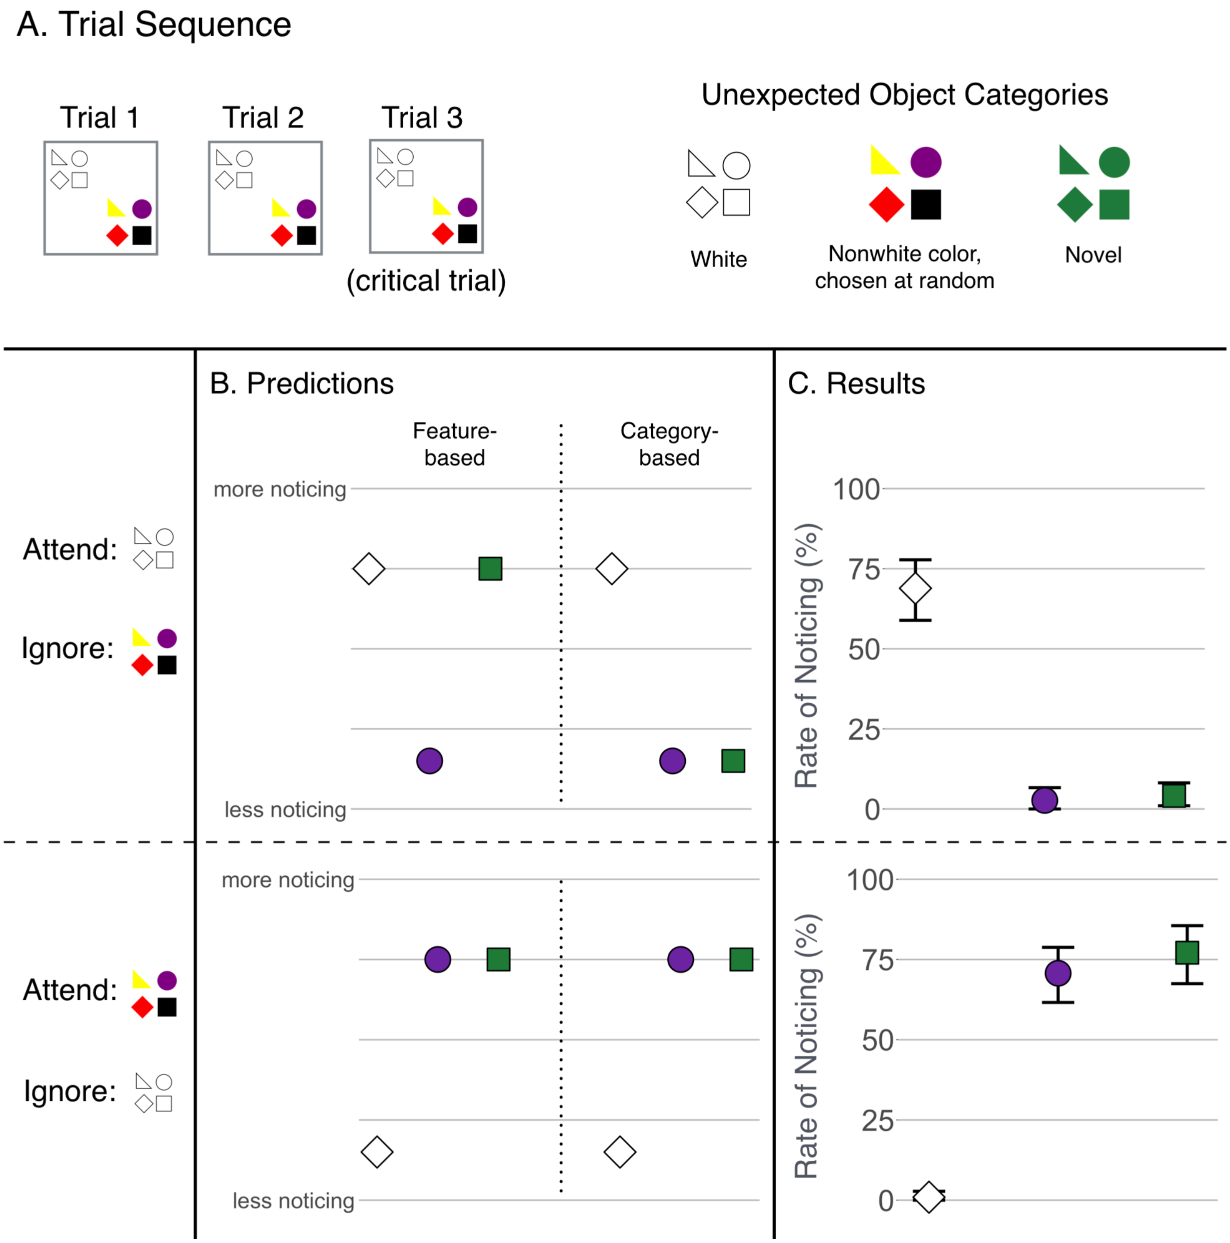

Figure 2.1. Trial sequence, predictions, and results for Experiment 1. **A.** A schematic of the objects in each trial. All four colors appeared on each trial alongside the white objects. Subjects attended either to white or nonwhite shapes. The unexpected object was a randomly chosen shape and could be white, one of the display colors (red, black, purple, or yellow, chosen at random), or a new color (green). **B.** If objects are sorted into attention sets on the basis of their features (the "feature-based" hypothesis), then a novel object should stand out in the display when ignoring colors. Conversely, if the objects are separated on the basis of categories (the "category-based" hypothesis), a nonwhite novel object should be categorized into the same set as the other nonwhite objects in the display. **C.** Noticing rates for unexpected objects when attending colors (top) and ignoring colors (bottom); error bars represent 95% bootstrapped confidence intervals. Unexpected objects that matched a display color were collapsed into a single group, represented by the purple circle. Novel green objects were noticed at virtually identical rates to an unexpected object that matched another color in the display when subjects ignored colors, suggesting a category-based attention set. All plots were generated with the ggplot2 package for R (Wickham, 2009).

or nonwhite shapes bounced against the edges of the rectangle, and to disregard the bounces of the other group, all while keeping their eyes focused on a small blue fixation square centered in the window.

Both sets of objects consisted of a square (44 x 44 pixels), a diamond (identical to the square, rotated 45 degrees), a triangle (50 pixel base, 50 pixel height), and a circle (22 pixel radius). For one set, all four shapes were white (#FFFFFF). For the other set, each shape was randomly assigned (without replacement) to be red (#E41A1A), yellow (#E8F212), purple (#6E24A5), or black (#000000) at the start of the experiment (see Figure 2.1, column A).

On each trial, these eight objects moved around inside the blue (#58ACFA) frame (666 x 546 pixels) for 17 seconds. Each object moved independently with a velocity that could vary between 54 and 108 pixels per second, reversing direction when it came into contact with an edge of the frame. Objects occluded each other when they crossed paths, but always remained behind the fixation square (11 x 11 pixels, #0000FF). On average, the set of four objects bounced a total of 28.6 times (*SD* = 2.1). After the trial ended, subjects were instructed to enter their bounce counts in a text box that restricted their response to integers between 0 and 99.

Subjects completed two non-critical trials in which they counted bounces and entered their responses afterward. On the third, critical trial, an unexpected object entered the display from the right edge after 5 seconds, moved horizontally along the midline from right to left, passed behind the fixation cross, and exited on the left edge of the rectangle 6750 ms after it first appeared. This object was randomly selected to be one of the four shapes in the display, and could be either white, one of the non-white colors (red, black, purple, or yellow), or green (#1B7E39). After the trial ended, subjects entered their count as usual, but then were asked to report whether they had noticed “anything extra on the last trial that did not appear on the previous trials” (“yes” or “no”). Regardless of their response, they were asked about the color (“black,” “yellow,” “orange,” “red,” “green,” “blue,” “purple,” “white,” or “none of these”) and shape (“square,” “cross,” “circle,” “triangle,” “diamond,” or “none of these”) of the additional object.

Next, subjects reported their age, gender, country of residence, use of vision correction, status of their color vision, identification of Ishihara Plate 9, whether they encountered any issues with the display of the experiment, and whether they had any previous exposure to inattentional blindness tasks. After completing these questions, subjects received a completion code to enter on Mechanical Turk to receive payment.

Results and Discussion

Using our preregistered exclusion criteria, we excluded data from 448 subjects (44.8% of our sample), leaving 553 in the final analysis (see Table 1.1 for the number of subjects assigned to each condition). According to our preregistered criteria, subjects were counted as having noticed the unexpected object if they they reported noticing an extra object on the critical trial and correctly reported its shape, color, or both.

Regardless of whether the category-based hypothesis or the feature-based hypothesis is correct, white unexpected objects should be noticed more often when attending to white than when ignoring white, and colored unexpected objects that match colors already in the display should be noticed more when attending to colors than when attending to white. As predicted, the difference in noticing rates between the attend-white and attend-nonwhite conditions was large and positive for white unexpected objects (68.9% versus 0.9% for a difference of 68%, 95%

bootstrapped CI: [58.9, 75]) and equally large and negative for colored unexpected objects (2.7% versus 70.7% for a difference of -68%, 95% bootstrapped CI: [-72.2, -60.6]; see Figure 2.1, column C).

The feature-based hypothesis makes a different prediction about the difference in noticing rates for a new, green object than does the category-based hypothesis. If the individual features of each object are used to form the attention set, then a green object should be noticed at a high rate regardless of whether people are attending to white or colored shapes (no difference in noticing between conditions; Figure 2.1, Panel B, “Feature-based”). Conversely, if attention sets are category-based, then the green object should be grouped into the category-based attention set for the nonwhite objects; It should be noticed rarely when attending white and ignoring colors, but frequently when attending colors and ignoring white (a large negative difference; Figure 2.1, Panel B, “Category based”). The results match the predictions of the category-based model of attention sets: Green unexpected objects were noticed rarely when ignoring colors, but noticed frequently when attending them (4.1% versus 77.1% for a difference of -73%, 95% bootstrapped CI: [-77%, -66%]; see Figure 2.1, column C).

Apparently, the attention set formed in this task is category-based rather than feature-based. People appear to suppress anything matching the category “colored” or “non-white” rather than specifically suppressing black, red, yellow, and purple. Such category-based attention sets have been shown before with semantic categories (Most, 2013), but this experiment suggests that even simple visual stimuli are coarsely categorized when forming attention sets.

EXPERIMENT 2

The category-based suppression observed in Experiment 1 might not apply to all cases of inattentional blindness. Perhaps people form attention sets flexibly, using category-based selection when conditions allow it and feature-based selection in other contexts. Asking people to ignore or attend to four colors at once may have taxed working memory, so forming a category-based attention set to cover all nonwhite objects was the most efficient way to perform the task. If that load were reduced, would people still form a category-based attention set? In other tasks, reducing perceptual load increases awareness of irrelevant stimuli (Cartwright-Finch & Lavie, 2007), and the same might hold for sustained inattentional blindness.

In Experiment 2, we reduced the load on working memory while ensuring that the absolute color of the non-white stimuli remained irrelevant. On each trial, rather than presenting four different colors simultaneously, the non-white set was composed of just one color. That color changed on each trial (e.g., white versus purple, then white versus red, then white versus yellow). On each trial, then, subjects still attend to white or color, but they have minimal memory load—they can ignore the category of "color," or selectively ignore the single color presented on the current trial.

If working memory demands alone drove the pattern of results observed in Experiment 1, then when people attend white and ignore colors, we should only observe suppression for an unexpected object that matches the current color on that trial; previous colors and novel colors (i.e., green) should escape suppression because they differ from the ignored color (Figure 2.2, Panel B, "Working Memory"). If attention sets are category-based even when the memory load is minimal, then any colorful unexpected object should be suppressed, even when it differs from the color on that trial (Figure 2.2, Panel B, “Variability”).

Methods

A version of the experiment that does not collect data but is otherwise identical to the one used in the experiment can be found here: http://simonslab.com/mot/color_demo.html. Except where noted, the methods were identical to those of Experiment 1. We recruited 1002 subjects through Mechanical Turk, again with the goal of 100 subjects per condition after exclusions.

The non-white objects on each trial all shared a single color, but the color was different on each trial (red, yellow, and purple, ordered randomly for each participant; see Figure 2.2, column A). The unexpected object was a randomly selected shape, and could be either one of the colors used for the display shapes in the non-critical trials, the same color as the non-white objects on the critical trial, or a new color (green). The same response options were available for the appearance of the unexpected object except “orange,” which was dropped from the possible colors.

Results and Discussion

After applying the same exclusion criteria as in Experiment 1 (eliminating data from 467 subjects—46.6% of our sample), our analyses included 535 subjects (see Table 2.1 for a breakdown of condition assignment).

Consistent with the the results of Experiment 1 and inconsistent with the memory load account, all colors were noticed less often when attending white and ignoring a single color than when attending to a single color and ignoring white. Even unexpected objects with colors that were unique in the display on the critical trial were missed. The difference in noticing an

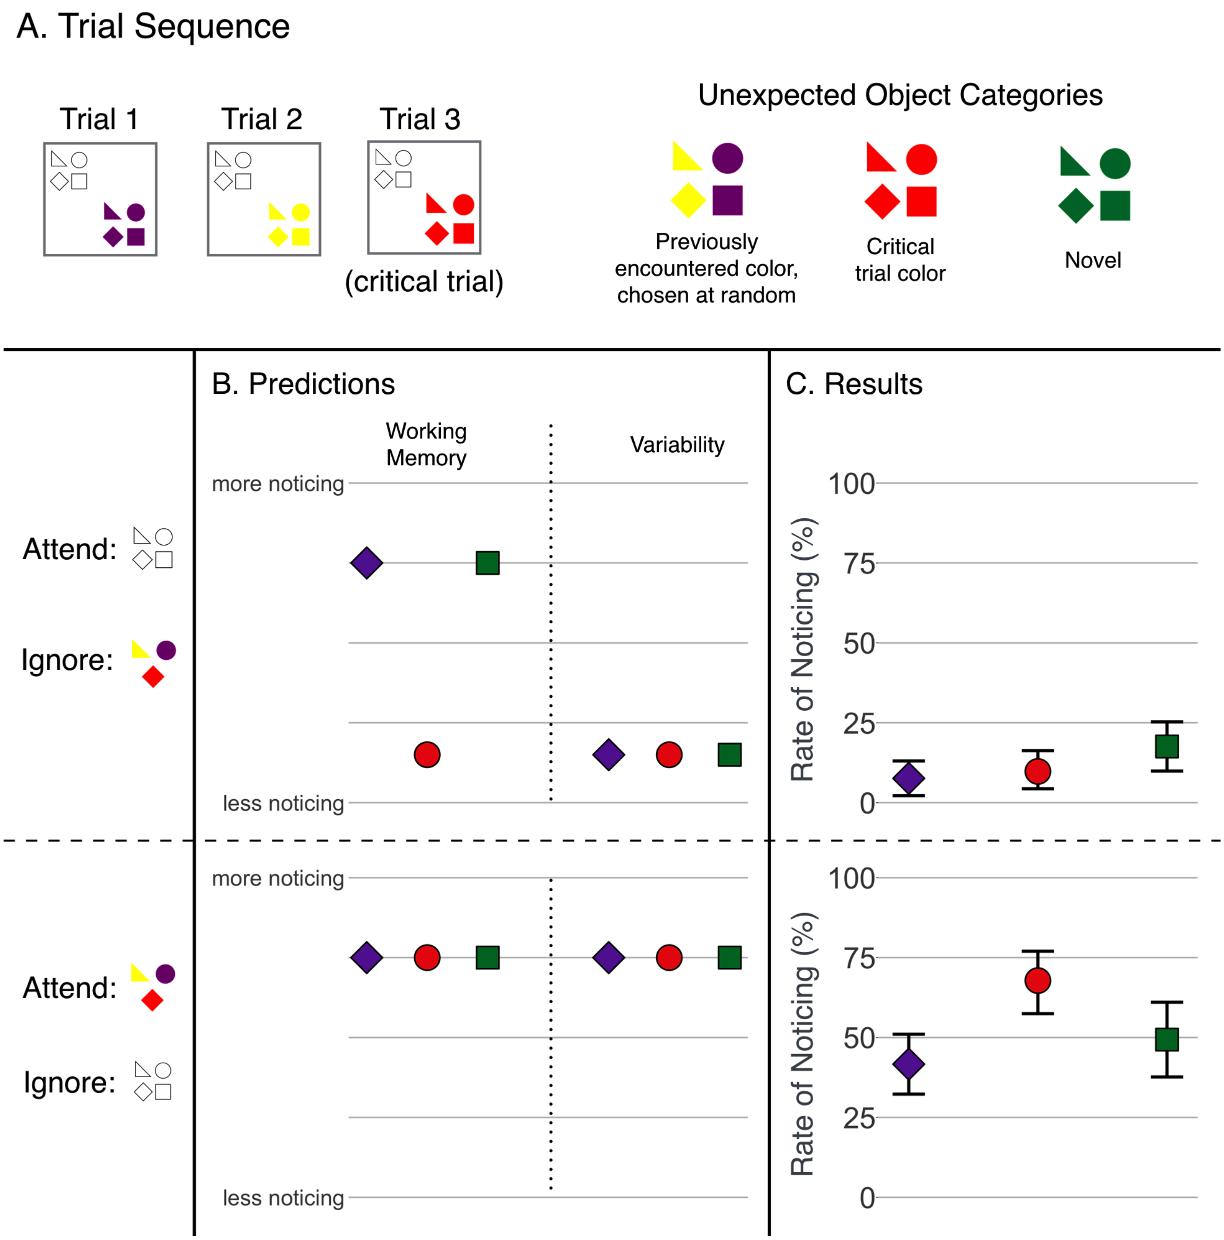

Figure 2.2. Trial sequence, predictions, and results for Experiment 2. **A.** A schematic of the objects in each trial. The nonwhite color was consistent within a trial, with only one color appearing alongside white, but this color changed on each trial. Subjects either attended white or attended the color. The unexpected object was a randomly chosen shape and could match a previously encountered color (in this example sequence, this would be purple or yellow), match the color on the critical trial (e.g. red), or be a new color (green). **B.** If broad filters are only a result of working memory load, then when working memory load is relieved, only a feature-based similarity effect is expected, with all colors standing out except the one currently ignored when ignoring color. If filters are always broad in the presence of variability, all colors should be suppressed when ignoring any color, and all noticed at high rates when attending any color. **C.** Results. All unexpected objects that matched a previously encountered color are collapsed, represented by the purple circle; those matching the color on the critical trial are collapsed and represented by the red square. Error bars are 95% bootstrapped confidence intervals. All colors (previously encountered, matching the critical trial, and new) were noticed at low rates when ignoring colors, and noticed more frequently when attending colors (although the current match was noticed most often).

unexpected colored object when ignoring a color versus when attending to a color was negative in all cases (unexpected objects that matched a color encountered on an earlier trial: 7.6% versus 41.7% for a difference of -34.1%, 95% bootstrapped CI: [-38, -29.1]; unexpected objects that matched the current color: 9.8% versus 67.8% for a difference of -58%, 95% bootstrapped CI: [-60.8, -53.1]; unexpected objects in a novel color: 17.6% versus 49.4% for a difference of -31.8%, 95% bootstrapped CI: [-35.6, -28]; see Figure 2.2, column C).

Despite having to attend to and ignore just one color on each trial, subjects formed an attention set for the entire category of "color," effectively suppressing detection of unexpected colors that were unique to the display on that trial. Apparently, the selective suppression of objects from the broad category of "color" in Experiment 1 was not because working memory was overloaded, but because having to ignore a variable feature led people to establish a category-based filter.

When people attended to that variable feature, however, the pattern was not as consistent; an unexpected object that matched the currently-attended color was noticed more often than either a previously-attended color or a novel color. While people effectively suppressed a variable category, they seemed less likely to enhance one.

EXPERIMENT 3

The results of Experiment 2 present an intriguing possibility: When people ignore a varying set of objects and attend to something constant, any object that differs from the attended one is suppressed. Neither novelty nor task-irrelevance rescues these objects. Conversely, when people attend to a varying set of objects, everything is noticed more often, although even novel objects are still noticed less often than objects that perfectly match what the currently attended feature.

In Experiment 2, the constancy versus variability manipulation was confounded with pre-existing natural categories: “chromatic” and “achromatic” objects. The varying objects were also the colorful ones, and the constant objects were always white. Consequently, the different pattern of noticing when the attended and ignored items are variable might be explained by a difference in how well people can ignore or attend to “chromatic” objects as a category. If so, the pattern in Experiment 2 would be consistent with a similarity effect: chromatic objects go unnoticed when ignored but are detected when attended.

To test whether natural chromaticity categories drove the effect in Experiment 2, we replaced white objects with pink ones so that there were no “achromatic” objects, and let the color of the constant objects be a randomly selected color for each subject instead of fixed for all subjects. If the results of Experiment 2 were due to the use of “chromatic” as a category, then we should not replicate the pattern in Experiment 3 (Figure 2.3, Panel B, “Chromaticity”). However, if the results of Experiment 2 were due to the ease with which people enhance or suppress heterogeneous categories, then Experiment 3 should replicate the pattern observed in Experiment 2; noticing rates for all colors of unexpected object should be lower when ignoring the varying color objects than when ignoring the constantly colored one (Figure 2.3, Panel B, “Replicate Ex. 2”).

Methods

A demo version of the experiment may be found here: http://simonslab.com/mot/scramble_demo.html. Except where noted, methods are identical to those of Experiments 1 and

2. Because the effects we were aiming to replicate from Experiment 2 were large, for this experiment we aimed to recruit 50 per cell for a total of 300 after exclusions. We recruited 609 subjects through Mechanical Turk.

The experiment procedure was identical to Experiment 2, except that the white objects were replaced with pink ones (#FFC0CB).Rather than a single constant color (white in Experiments 1 and 2), each subject was randomly assigned a “constant” color that appeared as one of the sets on every trial. The other set of objects all shared the same color, selected from the non-constant ones for that participant (see Figure 2.3, column A). Subjects were assigned to attend either to the constant color or to the color that changed from trial to trial. The unexpected object could be entirely novel (green), match the non-constant color from a previous trial, or match the non-constant color on the critical trial.

Results and Discussion

234 subjects were excluded prior to analysis (38% of our sample) according to the same exclusion criteria used in Experiments 1 and 2, leaving 373 subjects (see Table 1.1 for the numbers assigned to each condition).

Unlike in previous experiments, here subjects could not distinguish the sets of objects based merely on the presence or absence of color; all objects were chromatic, and there was no systematic relationship between the colors of the objects. The only way to separate the objects without using individual features is to use categories such as “the constant color” (e.g. “yellow”) and “the other colors” (e.g. “not yellow”).

Despite the absence of a pre-existing category, Experiment 3 replicated the results of Experiment 2 when subjects ignored the varying color (see Figure 2.3, column C). In all

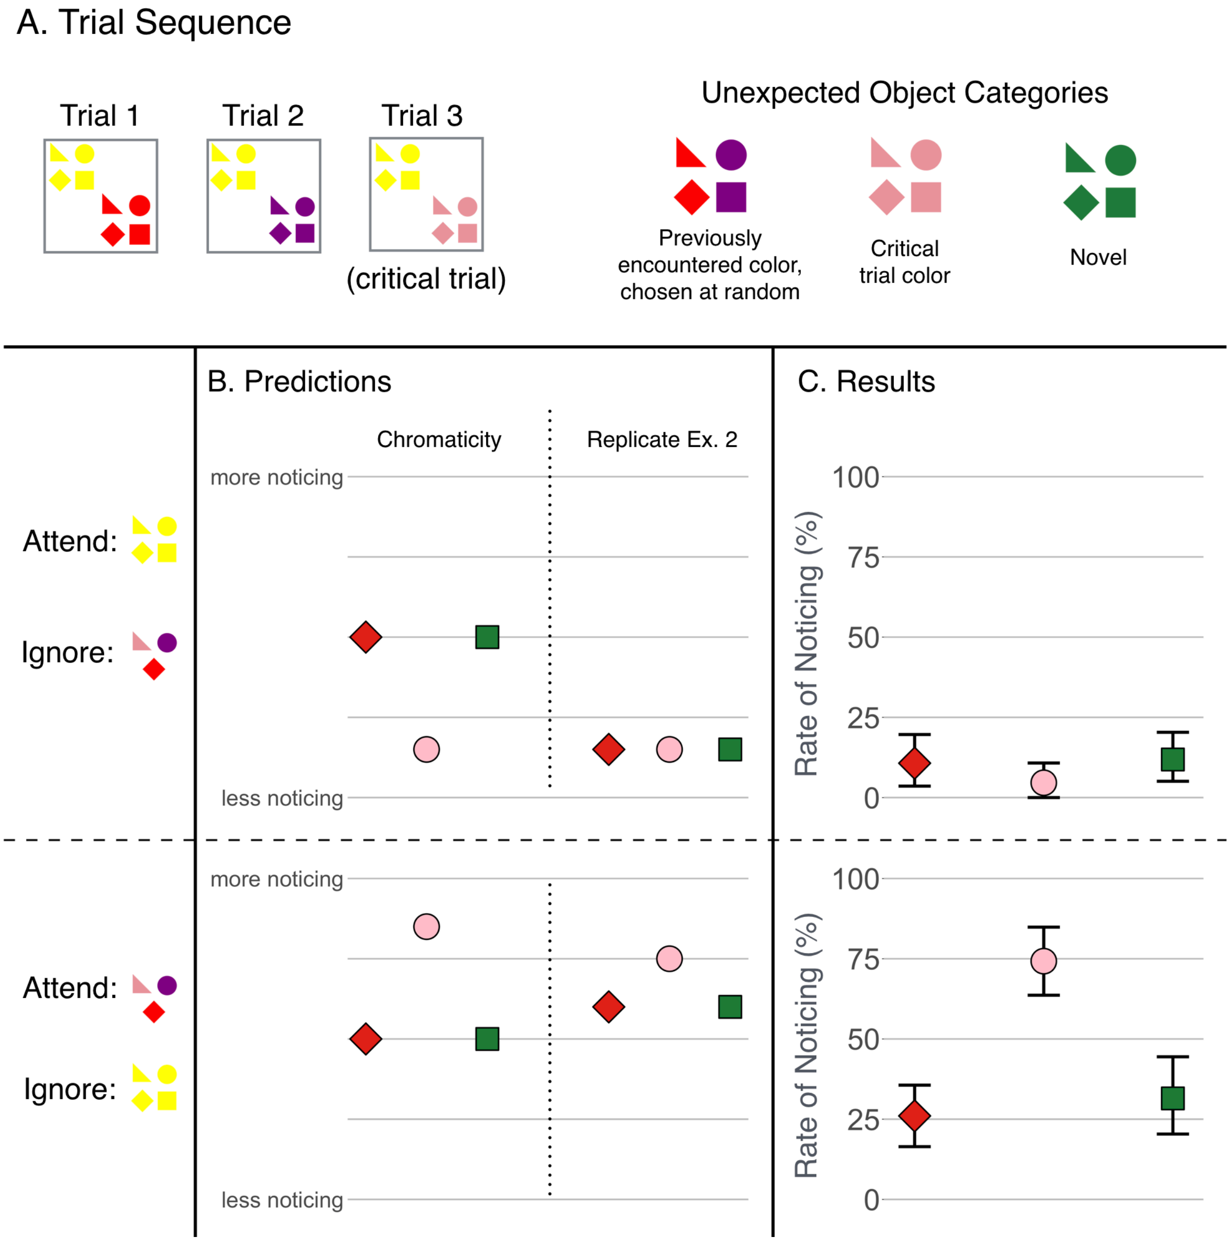

Figure 2.3. Trial sequence, predictions, and results for Experiment 3. **A.** A schematic of the objects in each trial. Rather than using white as the constant color, it was replaced with pink. A color was chosen at random to be constant for each subject. The other three colors served as the varying colors from trial to trial. Subjects either attended the constant color (yellow, in this sample sequence) or the varying color (red, purple, and pink). The unexpected object was a randomly chosen shape, and could match a previous color from the changing set (red or purple), match the current varying color (pink), or be a new color (green). **B.** If we replace white with any unchanging color, then broad variance-based filters should respond the same as in Experiment 2. Otherwise we might expect a pattern where we see strong similarity effects for the current color, and middling noticing for the other colors. **C.** Unexpected objects in a previously encountered color are collapsed and represented by the red diamond; those representing the color on the critical trial are represented by the pink circle. Error bars represent 95% bootstrapped confidence intervals. The same pattern of results emerges as in Experiment 2. When ignoring the changing color (top), all colors are suppressed, not just the one being actively ignored. When attending the changing color, there were even greater differences between noticing rates for an unexpected object that matched the currently attended color versus other colors than in Experiment 2 (bottom).

conditions, the difference between noticing rates when ignoring the varying colors versus attending to them was negative; subjects who ignored varying colors were always more likely to miss a colored unexpected object (unexpected objects that matched a varying color from a previous trial, 10.7% versus 26% for a difference of -15.3%, 95% CI: [-17.8, -12.5]; an unexpected object that matched the current color from the variable set, 4.6% versus 74.2% for a difference of -69.6%, 95% CI: [-74.1, 63.6]; an unexpected object in a novel color, 11.9% versus 31.5% for a difference of -19.6%, 95% CI: [-24, 15.1]). It appears that people can establish attention sets solely on the basis of constancy and variability.

One interesting pattern, present in Experiment 2, was amplified in Experiment 3. When subjects were attending to the variable colors, they noticed all unexpected objects more often than when ignoring the variable colors. However, whereas the color they were attending was noticed at fairly high rates (74%), the other colors were noticed considerably less often (26% for the previously attended colors and 32% for the new color). The effects of suppression were consistent across experiments: When people ignore varying colors, whether they vary within a trial (Experiment 1) or across trials (Experiments 2 and 3), they suppress all non-attended colors. In contrast, Experiments 2 and 3 revealed that when people attend to varying colors, noticing is only enhanced for the currently attended color. There is no increase in noticing for previously attended colors or for novel colors.

GENERAL DISCUSSION

Across three experiments, people formed attention sets that suppressed variations in color, leading to reduced noticing of unexpected objects of any color, irrespective of task-relevance or novelty. When one set of objects was constant and the other was heterogeneous within a trial or homogenous but variable across trials, subjects established an attention set that suppressed detection of other colorful shapes, even unique ones. Although the suppressive effect of ignoring a set of shapes seems to be broad in scope, the enhancement of attended features appears to be narrower. People tend to notice unexpected objects that match the current attended feature more than unique colors or colors that match previously attended features.

In Experiment 1 people either attended white and ignored a set of colorful shapes, or vice-versa. The two models for attention sets—category-based and feature-based—make different predictions for noticing of an unexpected, novelly colored object when attending to white objects and ignoring colored objects. Consistent with the category-based selection model, noticing of a unique green shape was just as suppressed as noticing of other colors in the display, despite being featurally distinct from the ignored set.

Experiment 2 confirmed that the broad suppression was due to variability of the color rather than to working memory load: Even when the colors within a trial were homogenous, variation in the color across trials affected the attention set in a category-based way. When subjects ignored the varying color set they missed not only unexpected objects that matched the color on the critical trial, but also colors they had ignored before and a completely novel color not yet encountered in the display. Every color was filtered, even if it was no longer task relevant or had never been encountered before.

In contrast, when attending to the varying color and ignoring white objects, an unexpected object matching the color on that trial was noticed more often than unexpected objects matching previously encountered colors or a novel color. These differences increased in Experiment 3, when we replaced white with pink and designated one color at random to remain fixed throughout the experiment. Experiment 3 also confirmed that selective ignoring operates in a category-based way in the face of variability, even when eliminating a possible simple feature dimension (chromaticity) as a way to establish an attention set.

This asymmetry in effects for selective attention and selective ignoring is suggestive. In previous inattentional blindness work, the pattern of results often reverses when the attended and ignored items are swapped. For example, when you attend black and ignore white, you see black often but rarely catch white; when you ignore black and attend white, the data reverse (Most et al., 2001). In this case, however, we observe two different data patterns when subjects attend to variability versus ignore it. When subjects ignore variability, people apparently suppress ignored objets in a category-based way, suppressing objects with features that are not part of the current display. In contrast, when attending to variability, selection appears to be more feature-based. Unexpected objects that perfectly match the attended ones are noticed at a high rate, but other objects are noticed less frequently.

CONSTRAINTS ON GENERALITY

All three experiments establishing filtering using color as the critical feature. Given that similarity effects in inattentional blindness research have been studied with a wide variety of stimuli (Simons and Chabris, 1999; Most et al., 2001; Most et al., 2005), we expect this effect to generalize to other kinds of simple objects, provided that a “constant” category is pitted against a “variable” category. If the objects must be separated along a single feature dimension, such as shape, luminance, or color, and one set contains members that are heterogeneous with respect to this critical feature, we would expect to observe the same general pattern (broad suppression when ignoring the variable group, and narrow selection when attending to it). The effects may even be stronger the more difficult it is to segment the objects, based on how noticing rates changed between Experiments 2 and 3; white was apparently easier to ignore than another color, so noticing rates were higher when ignoring white than when ignoring a randomly selected, unchanging color. However, these results might not generalize to more complex objects or richer stimulus categories that vary along more than one simple feature.

Although it is possible that the pattern we observed would vary with different task instructions, we expect that they would be robust to such variations given that people tend to partition sets of objects in simplest way they can, even when they are instructed to use a different feature; for example, when given white diamonds and black squares and told to attend squares and ignore diamonds, noticing of unexpected objects suggests that people use luminance instead (Aimola Davies, Waterman, White, & Davies, 2013).

We used a relatively diverse online sample, and given that inattentional blindness studies have shown effects of similarity in both online studies and in laboratory settings and to subjects of varying ages, we expect our pattern of results would generalize to any population of adult subjects who meet our inclusion criteria (although absolute noticing levels might vary across samples).

CONCLUSION

Previous studies of the contribution of attention sets to inattentional blindness appeared to provide evidence for sets based on features, relations, and even semantic categories. However, these options are indistinguishable in most previous work because the stimuli in the display could be distinguished in multiple ways. For example, black and white shapes differ in absolute RGB value, relative luminance, color class, and so on. Unexpected objects also fall nearer to one of the sets of objects according to multiple criteria: a “darker” object is also closer to black in RGB space, meaning attention sets based on relations and features make the same predictions for noticing rates.

In our experiments, the display objects were heterogeneous with respect to color, the critical feature. While one set of objects all shared the same color, the other set varied, either within trials or across them. With variation in color, the different types of attention sets made different predictions for the patterns of noticing for unexpected objects; a feature-based approach predicted a pattern of noticing completely distinct from that predicted by a coarser, category-based approach.

Across three studies, subjects formed broad, category-level attention sets that seem consistent with two categories: “attended objects” and “everything else.” The “everything else” category contains not just distractors that need to be immediately ignored, but also stimuli that have been encountered before and even completely novel stimuli. Indeed, the only type of unexpected object that was consistently noticed at a high rate was one that exactly matched the currently attended objects. The objects we must immediately contend with gain access to our awareness; all others, old and new, are suppressed.

CHAPTER 3: UNDIRECTED SPATIAL ALLOCATION

This chapter describes previously published experiments.[2]

Subjects in visual attention experiments rarely get to set their own priorities. They are explicitly instructed where and to what they must attend, advised what information within the display will be relevant or irrelevant to the task, and may be asked not to move their eyes during the task. This allows experimenters a fine degree of control over when and where attention is allocated, at the cost of allowing subjects to select their own approach to the task.

When displays and tasks are not completely constrained, certain characteristics of selective attention make themselves evident. For example, if the to-be-attended and to-be-ignored objects in a display can be distinguished from one another by either shape or luminance, subjects appear to select based on luminance irrespective of the feature they are explicitly instructed to use (Aimola Davies, Waterman, White, & Davies, 2013), likely because luminance or color is easier to attend to (Treisman & Souther, 1985; Wolfe, 1994). Feature-based attention also does not appear to operate across the entire visual field if the task does not require that it do so, but rather only operates in spatially relevant areas (Stothart, Simons, Boot, & Wright, 2019). This suggests that selective attention prioritizes the features that are easiest to operate on by default, and only maintains attention sets over relevant regions of space.

How do we set our attentional priorities when there are no requirements to attend to or ignore anything in particular, and when attention is not directed to any particular location? To what degree does a particular environment influence the allocation of attention? Inattentional

---

[2] Wood, K., & Simons, D. J. (2019). The spatial allocation of attention in an interactive environment. *Cognitive research: principles and implications*, *4*(1), 13. doi:10.1186/s41235-019-0164-5. Published and reproduced under a CC-BY license (http://creativecommons.org/licenses/by/4.0/).

blindness paradigms are particularly well-suited to address these questions. Because the critical stimulus is unrelated to the primary task, it does not affect how subjects allocate their attention before it appears. It nevertheless acts as a measure of the degree to which attention is focused on a particular feature or region of space.

When constraints are removed from a task, subjects appear to take the path of least resistance. They partition display objects into groups according to the easiest feature, and they do not maintain feature-based attention sets in irrelevant parts of the display. What determines the focus of attention when the task is further unstructured, and subjects do not need to deliberately attend to or ignore any aspect of the display? Can the requirements and structure of the task alone drive an attention set for particular features or regions of space, encouraging selection within the attention set and inhibition of anything falling outside of it?

In the following series of experiments, subjects play a simple video game as their primary task. The game has a structured environment and obstacles that follow certain predictable patterns. Subjects are instructed how to play the game, but receive no other guidance as to strategy or where they should be attending. Using unexpected objects that appear while they play, it is possible to unobtrusively map out where they allocate attention and how this relates to the structure of the task they must complete. These experiments demonstrate that subjects concentrate their attention narrowly in the most task-relevant area of the display and disregard the rest, suggesting that attention automatically conforms to the constraints of the task in which it is deployed.

INTRODUCTION

The same place can seem entirely different depending on how we move through it. When walking to our regular coffee shop, we concern ourselves with navigating around other pedestrians on the sidewalk and checking that streets are safe to cross, paying little mind to the cars passing by. Driving there places different demands on attention; we would focus on the cars, crosswalks, traffic signals, and open parking spaces, paying little heed to pedestrians on the sidewalk. When a task requires us to focus our attention on a particular region of space, we appear to ignore or filter out task-irrelevant areas.

We engage in this filtering of regions outside our focus even when performing straightforward tasks with simple displays that require no walking, or even eye movements. When subjects focused on a cross in the center of an otherwise empty display and judged which of its arms was longer, they were less likely to notice a new, unexpected object the further it appeared from the cross (Newby & Rock, 1998). Similarly, when counting how many times a subset of the moving objects in a display crossed a horizontal line bisecting the display, subjects were increasingly less likely to notice unexpected objects the further they were from the line (Most, Simons, Scholl, & Chabris, 2000; Stothart, Boot, & Simons, 2015).

Inattentional blindness methods are especially well-suited to studying the effects of proximity to the focus of attention. Because the critical object appears unexpectedly, subjects have no reason to divert attention from their primary task, or to attend to or ignore objects they might not otherwise. Other studies probing the spatial characteristics of the “attentional spotlight” do not have this advantage, instead often deliberately interfering with the primary task by cuing movement of attention away from the stimulus (e.g. Posner, Snyder, & Davidson, 1980) or employing highly confusable distractor stimuli near the focus of attention (e.g. Müller, Mollenhaur, Rösler, & Kleinschmidt, 2005). However, for the inattentional blindness tasks used to study proximity to date, the spatial layouts of the displays were arbitrary and dictated by the task. Unlike the pedestrian strolling to the coffee shop, their actions do not guide attention in these tasks. There is nothing inherent to these displays that would naturally direct attention to a particular area, and subjects are not interacting with the displays themselves beyond making judgments or counting with their eyes fixed on one spot. The role of context and task is left an open question in these particular paradigms.

Tasks in which subjects interact with an environment in some way reveal an influence of this interaction on how and where they allocate our attention. When subjects are moving through a road-like setting, they show worse change detection performance while actively steering themselves compared to when they were “passengers” (Wallis & Bülthoff, 2000). However, the active “drivers,” while worse overall, detected changes better near the center of the road than changes farther from it. The demands of driving apparently narrowed the scope of attention to elements closer to the road. However, this task too relied on change detection as the primary task, which was unrelated to the act of navigating the environment.

Consistent with the effect of action demands on attention, both novice and expert drivers fixate on the road one to two seconds ahead, but the patterns of fixations vary depending on the kind of road (Underwood, Chapman, Brocklehurst, Underwood, & Crundall, 2003). On low-traffic rural roads, drivers tended to spend more time looking straight ahead. On roads with merges, they tended to check their mirrors more frequently. Although it is reasonable to assume that attention follows gaze with these drivers, and that they pay more attention to the road straight ahead when they do not have to execute any complicated maneuvers or respond to other vehicles, we cannot tell from observing patterns of eye movements alone where attention is allocated.

Studying the effects of how subjects interact with their environment on the allocation of attention requires a task with several properties: (a) unobtrusive measurement of attention; (b) sufficient freedom to make the actions seem natural; and (c) enough control to allow systematic measurement of where attention is allocated. We developed a simple road-crossing game in which subjects shuttle objects between safe zones (sidewalks), avoiding obstacles along the way and earning bonus points for speed. We use an inattentional blindness paradigm in which our primary measure is the likelihood of noticing an unexpected object as a function of its position in the display.

Across several experiments, we use this task to address a number of questions. Most importantly, how do the constraints of an environment influence the allocation of attention when all subjects need to do is interact naturally with it? Inattentional blindness tasks like ours are especially well-suited to address this question because they measure attention unobtrusively. Subjects do not have to split their attention between interacting with the display and performing an unrelated secondary task. Furthermore, because we measure attention using an unexpected object—rather than one that is always present but ignored, or a rare but not unexpected object—we can be confident that subjects are not deliberately allocating attention to the object or adopting a goal of detecting it.

Our specific implementation of an inattentional blindness task also allows us to ask whether various environmental constraints, such as the means by which subjects can travel and the behavior of hazards, influence attention and noticing. We can further examine whether the behavior of the unexpected object itself influences noticing beyond what we might predict based on the demands on attention induced by the task environment alone.

GENERAL METHODS

Subjects

The need for signed consent was waived by the University of Illinois Institutional Review Board due to the low-risk nature of the experiment. The subjects in all experiments were US-based workers recruited through Amazon's Mechanical Turk service. We used TurkGate (Goldin & Darlow, 2013) to screen out subjects who had previously participated in experiments from our lab based on their worker ID. Subjects were directed to an external website running the experiment in Javascript, and upon finishing the experiment, they received an entered a completion code to receive payment ($0.30) for the HIT ("Human Intelligence Task," the term for the jobs posted to MTurk).

Subjects were automatically recruited in batches of up to nine using the boto3 Mechanical Turk SDK (https://github.com/boto/boto3). When we passed the recruitment threshold for an experiment, recruitment stopped and no further HITs were posted.

Results from previous studies from our laboratory using similar recruiting methods suggest that we could expect to exclude approximately 30 - 40% of all data collected. We set recruitment thresholds expecting to be able to use approximately 60% of the data in our final analysis for Experiments 1 and 2; however, exclusion rates were lower than anticipated, and so in Experiments 3 and 4 we recruited expecting 80% usable data. Because Experiment 5 used an unexpected object that differed substantially from the other experiments, we piloted that task

Object	**Maximum Horizontal x Vertical size (in pixels)**	**Color(s)**	**Speed(s) (in pixels/ second)**
Roadway	600 x 500	Dark gray (#777777)	NA
Sidewalks	150 x 500	Medium gray (#C1C1C1)	NA
Seed Basket and barn	220 x 200	Cartoon image	NA
Pentagon avatar	40x40	Purple (#800080)	180
Pedestrian triangles	30 x 30	Blue (#0000FF)	60 or 120
Car discs	44 x 44 (radius = 22)	Red (#FF0000)	60, 120, 180, or 240
Seed disc	16 x 16 (radius = 8)	#FFD700, #F08080, #FFA07A, #20B2AA, or #87CEFA	NA

Table 3.1. Appearance and behavior details for the objects used in the game. These parameters were consistent across experiments.

prior to data collection with small groups of about 40 subjects each (both were intended to test the effectiveness of the procedures and not to estimate the effects of interest). The overall procedure in the pilot was identical to that of the main experiment. The first pilot used a slightly different version of the unexpected object. The second verified that a substantial update to the Chrome browser released just prior to launching the experiment did not cause an increase in self-reported technical issues for subjects. Based on the exclusion rates for those pilot subjects, we recruited for 70% usable data in Experiment 5.

Materials and procedure

All experiments and analyses were preregistered on the Open Science Framework (OSF; https://osf.io/brk6t/wiki/home/). Each experiment was preregistered separately, prior to data collection for that experiment. Anonymized data, all experimental materials, analysis scripts, and preregistrations for each experiment are available on OSF.

Prior to the experiment, subjects were shown an information screen that provided experimenter and IRB contact information. It explained that their responses would be anonymous, described how their data would be used, and noted that their participation was voluntary. They were then presented with an instruction screen explaining how to play the game, and after they clicked through it the game loaded and began to run. The play area consisted of a road, bordered on either side by sidewalks (for a screenshot of the game, see Figure 3.1A; for detailed parameters of the game objects, see Table 3.1—note that because subjects completed the experiment on their own devices, screen size and viewing distance could not be controlled, so all distances and object sizes are given in pixels[3]). Blue triangle "pedestrians"appeared once every 400 ms, starting off-screen either above or below the display at a random horizontal position within the bounds of the sidewalk, and traveled either top-to-bottom or bottom-to-top either quickly or slowly. There were up to 10 pedestrians across both sidewalks on screen at once. Red circle "cars" emerged from the top of the screen, traveling top-to-bottom at a randomly selected speed. Cars appeared continuously throughout the task and up to 10 could be on screen at once. On the right side of the play area was a barn, and on the left, a basket of seeds.

Subjects controlled their avatar with the arrow keys, and it could only move in one direction. For example, when crossing from left to right, the avatar could only move to the right. While a key was depressed, the avatar moved at a constant velocity with no acceleration. The subject's avatar started on the left side of the screen at a fixed vertical distance of 300 pixels from the top of the game area, pointing toward the seed basket; they could only move right-to-left until they touched the seed basket, at which point their avatar picked up a randomly-colored seed and reversed to point towards the right side of the play area. Subjects could then only move

---

[3] In Javascript, a pixel refers to a CSS pixel rather than a physical pixel. CSS pixels scale automatically to the density of the display device, such that a single CSS pixel is drawn with more physical pixels on a higher-density display.

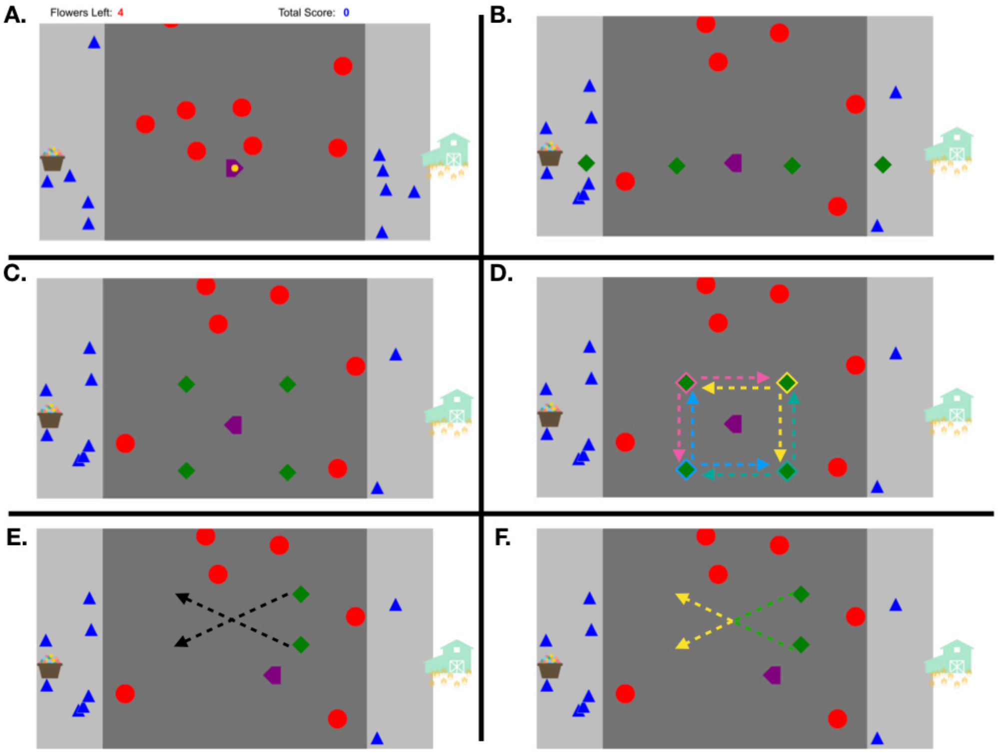

Figure 3.1. Experimental setup. **A.** A screenshot of the game, after a seed has been picked up. **B.** The four possible unexpected object positions in Experiment 1. **C.** The four possible unexpected object positions in Experiment 2. **D.** The four possible starting positions of the unexpected objects in Experiment 3. The dotted lines show the possible trajectories, and are color-coded to show which unexpected objects can take which trajectories. **E.** The two starting positions and corresponding trajectories for the unexpected objects in Experiment 4. **F.** The two starting positions and corresponding trajectories for the unexpected objects in Experiment 5. The color of the dotted lines corresponds to the color of the unexpected object at that point in its motion (the unexpected object could also start yellow and turn green).

left-to-right until they reached the barn. Subjects "planted a flower" when they carried a seed across the road and touched the barn on the opposite side, earning points equal to 50000 divided by the number of milliseconds they took to cross, or 1, whichever number was larger. Subjects had to plant 5 flowers in total to complete the task. If they contacted a car, their position was reset to the middle of the sidewalk from which they had begun that crossing. During the crossing either two or three crossings prior to the final one, an unexpected object appeared. The precise behavior of the unexpected object varied by experiment, but it was always a green (#008000) 40 by 40 pixel diamond and appeared abruptly in the display (i.e., a sudden onset). The primary question asked in all experiments was whether or not participants noticed this unexpected object as a function of its position and behavior.

After they finished the game, subjects were asked whether or not they noticed anything new that was not a game object. Regardless of their professed noticing, they then were asked: (1) whether the new object was moving, (2) in which direction the object was moving, (3) what color the object was (red, green, blue, purple, yellow, gray, black, white, or brown), and (4) what shape the object was (rectangle, triangle, diamond, circle, cross, T-shaped, L-shaped, B-shaped, or V-shaped). For experiments in which the unexpected object did not move, subjects were also asked about its location, either relative to the screen (right or left side) or relative to the subjects' avatar (above or below), depending on the experiment. Finally, subjects were asked to select their age range, gender, whether their vision needs correction and if they were wearing it during the experiment, the status of their color vision, the number contained in Ishihara Plate 9 (Ishihara, 1990), whether they had experienced any technical difficulties during the game, and whether they had prior experience with a similar inattentional blindness task. After submitting their final response, subjects were presented with a completion code and told to return to Mechanical Turk to enter the code and receive payment.

Analysis Software

All analyses were conducted in R version 3.5.1 (R Core Team, 2018) using packages ggplot2 version 3.0.0 (Wickham, 2016), stringr version 1.3.1 (Wickham, 2018), purrr version 0.2.5 (Henry & Wickham, 2018), tidyr version 0.8.1 (Wickham & Henry, 2018), and dplyr version 0.7.6 (Wickham, François, Henry, & Müller, 2018). Analysis scripts for each experiment were written and preregistered prior to data collection for that experiment and are available on OSF.

Analysis Procedure

For all analyses, we adopt an estimation-based approach. The target sample sizes we employed (100 per condition) allow us to estimate noticing rates within approximately ±10% across experiments. We report point estimates for noticing rates in all conditions, along with 95% bootstrapped confidence intervals calculated via the percentile method (Efron & Tibshirani, 1993). For comparisons of interest, we also calculate difference scores and their 95% bootstrapped confidence intervals. Due to the nature of our data, we elected to use bootstrapped confidence intervals rather than standard-error intervals because bootstrapped intervals do not exceed the bounds of the data and can be asymmetric.

Exclusion Criteria

Our preregistered criteria excluded data from subjects who reported being younger than 18 years old; who reported needing vision correction but not wearing it during the experiment; who reported any type of non-normal color vision; who incorrectly reported the number in the Ishihara plate; who reported that the game lagged, froze, or had some other problem; or who reported prior experience with inattentional blindness tasks. For a detailed breakdown of the exclusions in each experiment, see Table 3.2.

EXPERIMENT 1

If attention is guided by the demands of the environment in which it operates, it should be straightforward to predict where it will be allocated when the environment is constrained. In the game subjects play, the direction of travel is restricted to one direction—they can only move

Experiment	Excluded for age	Excluded for vision correction	Excluded for color vision	Excluded for Ishihara Plate	Excluded for technical issues	Excluded for prior IB experience	Total excluded
1	1	36	22	54	34	16	129
2	0	44	17	37	30	11	114
3	0	97	37	85	68	19	251
4	0	28	13	20	18	4	68
5	0	16	34	37	70	10	112

Table 3.2. A breakdown of the number of subjects excluded by each criterion in each experiment. A subjects could be excluded under multiple criteria, so the sum of the individual exclusions does not necessarily equal the total number of exclusions.

forward. Given that the risk of collision is always at or in advance of the subjects' current location, we might expect them to devote attention more to the region in front of their avatar than behind it. Similarly, because subjects must avoid colliding with objects while crossing the road, we might expect more attention directed to the regions of space nearest the subjects' avatar, in which the hazards pose the most threat, than to farther regions. As a result, we should expect more noticing for unexpected objects that appear near the subjects' avatar than far away, and more for objects appearing in front than behind the avatar. When collapsing across near and far conditions, a positive difference between noticing of unexpected objects appearing in front of subjects' avatar versus behind would suggest that more attention is allocated to the area in the direction of travel than to the inaccessible area behind the avatar. Collapsing across in front and behind, a positive difference in noticing of nearby versus far away objects would indicates that more attention is allocated nearby the avatar than farther away, possibly in order to successfully avoid obstacles.

Methods

A demonstration of the experiment, exactly as a subject would experience it but without any data collection, can be viewed at http://simonslab.com/game/crossing_demo.html.

Subjects. We aimed for usable data from 100 subjects per condition after exclusions (total target N=400). We set a recruitment target of 600 subjects and collected data from 634 in total.

Materials and procedure. Subjects were randomly assigned to one of four conditions, each corresponding to a possible unexpected object location relative to the player: near and in front, near and behind, far and in front, or far and behind (Figure 3.1B).

The unexpected object appeared either during the 7th crossing of the game, when subjects were carrying their 4th seed across the road, or on the 8th crossing, when they were returning to the seed basket to pick up the 5th and final seed (selected randomly). It was therefore random whether “in front” and “behind” corresponded to left or right. The unexpected object onset immediately when subjects crossed the midpoint of the game area (450 pixels from the edge) and remained visible for one second before disappearing. It appeared at the same vertical height as the subjects’ avatar (300 pixels from the top of the game area), either 113 pixels away horizontally in the near case or 338 pixels away in the far case. In the “in front” condition, the unexpected object appeared in the player’s path, and in the “behind” condition it appeared behind the player (i.e., in the direction their avatar could not travel). In the case of the near and in front condition, subjects could overlap with the unexpected object if they moved the entire time it was onscreen. The unexpected object occluded the avatar if they happened to intersect.

Results and discussion

Prior to analysis, we excluded data from 129 subjects (20.3% of our sample) according to the criteria in the General Methods. For our primary analysis, we coded subjects as having noticed the unexpected object if they correctly reported noticing something other than a game object, reported that it was not moving in response to both questions about the object's motion, and correctly reported which side of the screen (right or left) the object appeared on.

The noticing rates for the unexpected objects conform to the expected allocation of attention based on the demands of the display and the task. Subjects rarely noticed the "far behind" unexpected object, at 8.5% (95% CI: [4.6, 13.1]), but noticed the "near behind" object 47.5% (95% CI: [39.4, 56.2]) of the time. Noticing rates were higher for the unexpected objects that appeared in front of the subjects' avatar, with the "far in front" object noticed 38.1% (95% CI: [29.7, 46.6]) of the time and the "near in front" object noticed 69.2% (95% CI: [60.8, 77.5]) of the time.

An exploratory follow-up analysis examined whether the pattern of results differed if we counted a response as correct only following accurate identification of each of the unexpected object's features. We found no difference in the pattern of results regardless of the feature we required to be correctly identified (Figure 3.2).

People were more likely to notice unexpected objects that appeared near to their avatar than objects that appeared far away (a difference of 35.0 (95% CI: [27.1, 43.0]) percentage points, collapsing across in front and behind). There was a similar, 25.3 (95% CI: [16.9, 33.0]) percentage point advantage for objects that appear in front of the subjects' avatar versus behind it. It seems that people allocate their attention in response to the constraints of the environment,

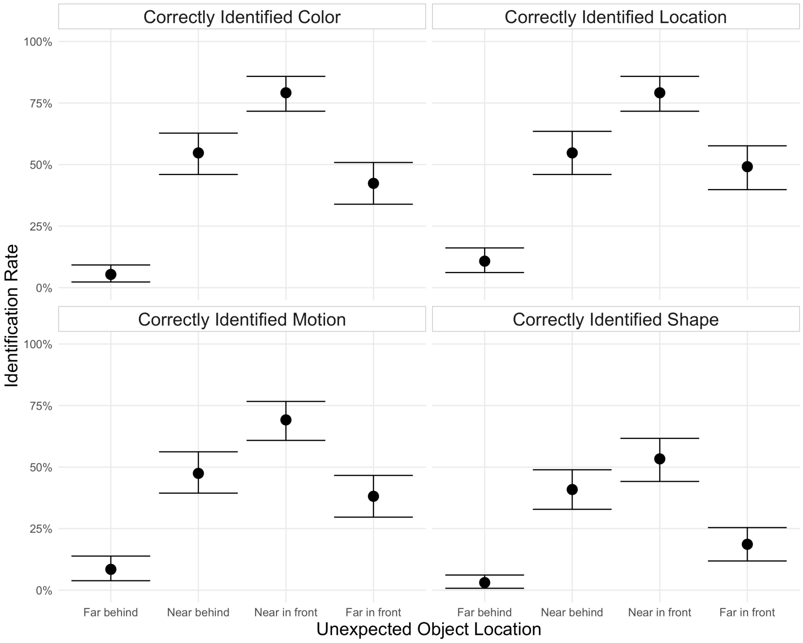

Figure 3.2. Rate at which subjects who reported seeing a new object successfully identified the unexpected object's features, broken down by each possible object position. Error bars are 95% bootstrapped confidence intervals. To be counted as correctly identifying a feature of the unexpected object, subjects first had to report noticing something new, and: for color, report that the new object was green; for location, report which side of the screen the object was on; for motion, report that the object was not moving; and for shape, report that it was a diamond.

with most of the attention directed near their avatar and in the direction of travel.

EXPERIMENT 2

The results of Experiment 1 confirmed that attention was allocated in response to environmental constraints. Subjects could only move in one direction, and unexpected objects were more likely to be noticed when they appeared in the path of the avatar's motion than when they appeared behind the avatar. Similarly, unexpected objects were more likely to be noticed when they appeared near the avatar than when far from it.

One question we might ask is whether this near-versus-far advantage results from the threat of collisions. Although unexpected objects whose features match those of threatening objects are not noticed more often than objects with features associated with neutral or rewarding objects in a game context (Stothart, Wright, Simons, & Boot, 2017), the hazards in our task might influence the spatial allocation of attention given their immediate consequences for action. If so, there should be differences in noticing rates for equidistant unexpected objects depending on where they appear relative to the subjects' avatar. An unexpected object that appears in front of and above the avatar, where there is the greatest danger of a collision (because the cars move from the top of the display to the bottom), ought to be noticed more often than an object the same distance away but beneath the player, where the risk of a collision has passed.

Experiment 2 uses the same methods as Experiment 1 to explore whether there is an above/below difference in noticing, similar to the near/far and in-front/behind differences observed in Experiment 1. When we collapse across the above/below conditions and examine the difference in noticing for the in front versus behind unexpected objects, we expect the same positive difference we observed in Experiment 1. Additionally, if more attention is directed to the high-risk areas above the avatar than to the areas below it, we expect a positive difference in noticing for unexpected objects appearing above versus below (collapsing across in front and behind conditions).

Methods

A demonstration of the task may be viewed at http://simonslab.com/game/updown_demo.html.

Subjects. We aimed for usable data from 100 subjects per condition, for a total of 400 subjects after exclusions. We recruited 540 subjects in total.

Materials and procedure. Experiment 2 used the display and task described in the General Methods and all details are identical to Experiment 1 except for the position of the unexpected object. In Experiment 2, the unexpected object could onset 122 pixels in front of or behind the player, and 122 pixels either above or below the player (Figure 3.1C) for a total of four conditions. If the avatar moved the entire time the unexpected object was onscreen, it would come level with an unexpected object that appeared in front of the avatar, but the avatar would not pass it before it offset. The post-game survey asked whether the unexpected object appeared above or below the subject's avatar (rather than what side of the screen it had appeared on as in Experiment 1).

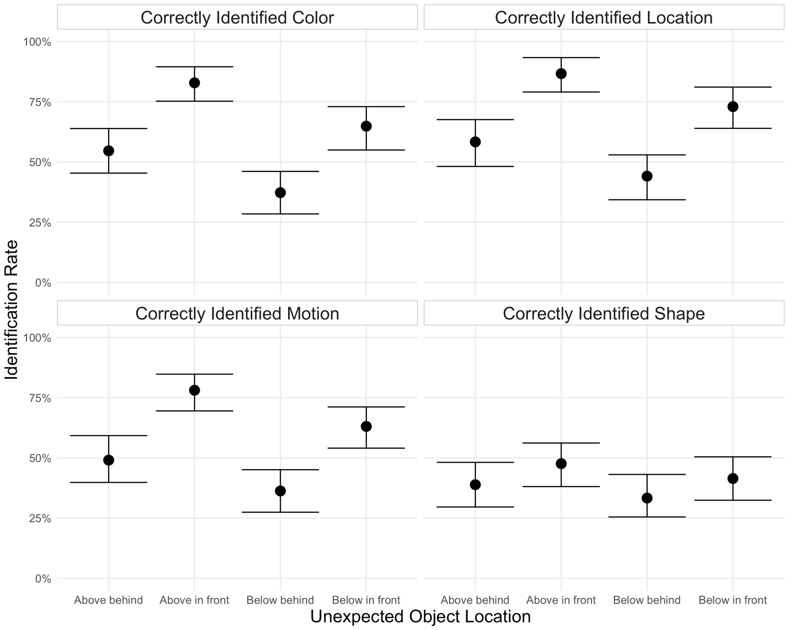

Figure 3.3. Rate at which subjects who reported seeing a new object correctly identified the unexpected object's features, broken down by each possible position. Error bars are 95% bootstrapped confidence intervals. To be counted as correctly identifying a feature of the unexpected object, subjects had to report noticing something new, and: for color, report that the new object was green; for location, report whether the object appeared above or below their avatar; for motion, report that the object was not moving; and for shape, report that it was a diamond.

Results and discussion

We excluded data from 114 subjects (21% of our sample) prior to analysis. As with Experiment 1, we classified subjects as having noticed the unexpected object if they reported noticing something new, said it was not moving, and correctly reported whether it had appeared above or below them. Among subjects reporting something new, the pattern of results was similar regardless of which feature we required to be correctly identified (Figure 3.3).

Overall, we found the same in front versus behind advantage as in Experiment 1, with unexpected objects appearing in the path of travel noticed 27.5 (95% CI: [18.6, 36.1]) percentage points more than objects appearing behind the avatar (collapsing across the above and below conditions). We also found a 13.2 percentage point advantage (95% CI: [3.8, 22.5]) for objects above versus objects below.

Noticing rates varied with unexpected object location. When the object appeared behind the subjects' avatar, it was noticed more when above the avatar (49.1%; 95% CI: [38.9, 59.3]) than below it (36%; 95% CI: [26.5, 46.1]). When the object appeared in front of the avatar, it was noticed more when above (78.1%; 95% CI: [70.5, 85.7]) than when below (63.1%; 95% CI: [55.0, 71.2]) the avatar.

The increased noticing for objects that appear above and in front of the subjects' avatar suggests that subjects are allocating their attention more heavily to areas in which they are at risk of colliding with a harmful object. Indeed, the "above and in front" unexpected objects had the highest noticing rate of any unexpected object in Experiment 1 or 2, even more so than objects that appeared directly in the path of travel. Subjects seem to be sensitive to the demands of the environment necessary for completing their task and they direct their attention accordingly.

EXPERIMENT 3

In Experiments 1 and 2, when participants performed a dynamic, goal-directed task in which they navigated an avatar through an obstacle-filled display, they monitored the space in front of their avatar more than the space behind it, the space above more than the space below, and nearby locations more than far away ones. Where an unexpected object appears relative to a subject's avatar has a substantial impact on its likelihood of being noticed.

The unexpected objects in Experiments 1 and 2 were all static and occupied the same region of space the entire time they were on screen. These static objects allow for a measure of the "attention spotlight" (Posner et al., 1980), but they do not allow an assessment of the dynamics of attention over time. In particular, the unexpected objects remain stationary while the avatar—and, presumably, the focus of attention—moves, changing the position of the objects to relative to attentionally relevant areas over time. Static objects do not provide a clear understanding of how objects moving in and out of the attended region interact with attention. Does the distribution of attention act only on space, so that if an object travels into a region of greater attentional relevance, it will be noticed more often, regardless of where it originated? Or does the distribution of attention apply not just to the space, but to all of the objects contained within it? That is, will an object that originates in an attentionally irrelevant area be noticed less often, even when it travels into an area of greater attention?

Results from early selective looking studies suggest that an object is no more likely to be noticed by virtue of passing into an attended area. In a task requiring subjects to count basketball passes between dark-shirted players and ignore white-shirted ones, subjects failed to notice a woman with an umbrella walking through the video, even when playback was stopped at a moment when the woman appeared to be kicking the tracked basketball (Becklen & Cervone, 1983). Passes frequently went through the woman and she often overlapped with monitored players, but noticing rates never exceeded 35%. However, as with other dynamic inattentional blindness tasks, subjects in this task passively observed the display, and the requirement to monitor three players across the screen precluded the narrow spatial distribution of attention we observed in Experiments 1 and 2 using our game task. The motion of the unexpected object may have a greater impact on noticing in our framework.

Experiment 3 presented moving unexpected objects in the same road-crossing task to examine these questions. Experiments 1 and 2 revealed substantial differences in the likelihood of subjects detecting unexpected objects depending on where they appeared; Experiment 3 explored whether similar differences exist for objects that onset in relevant areas and offset in irrelevant ones (or vice-versa). Collapsing across the unexpected object's trajectory allows us to verify whether the overall advantage for unexpected objects appearing above versus below and in front versus behind still emerge. Collapsing across position, we can determine the difference in noticing rates for unexpected objects that start in an irrelevant area and move into a relevant one (or the reverse) for horizontally and vertically moving objects. A positive difference would suggest an advantage for objects that move into a relevant area, a negative difference would suggest an advantage for objects that start in a relevant area, and no difference would suggest that the type of motion does not have a substantial impact on noticing.

Method

A demonstration version of the task with no data collection may be viewed at simonslab.com/game/transit_demo.html.

Subjects. We recruited 1000 subjects to get 100 per condition for eight conditions. Subjects were recruited according to the procedure outlined in the General Methods, and we collected 1082 in total.

Materials and procedure. The gameplay aspect of the task was unchanged from the General Method; the only adjustment to the method concerned the unexpected object. The unexpected objects appeared in one of the four positions used in Experiment 2; 122 pixels above or below the center of the display, and 122 pixels above or behind the center of the display. However, rather than appearing when the player had crossed the halfway point of the display, they appeared when the player had traveled 360 pixels (90 pixels shy of the halfway point). The unexpected object appeared and began moving at 240 pixels per second, traveled 244 pixels in a particular direction, and was onscreen for 1016 milliseconds. Because the unexpected object moved slightly faster than the avatar and appeared when the avatar had not yet reached the midpoint of the screen, a horizontally-moving unexpected object would spend half of its time in front of the subjects' avatar and half behind (assuming the avatar moved continuously while the unexpected object was onscreen), and a vertically-moving unexpected object would spend half its time above the avatar and half below. Due to the positions and speeds of the objects, the unexpected objects always offset at least 60 pixels ahead of the avatar in the horizontal direction regardless of how much the avatar moved while the object was on screen.

The unexpected object could travel either horizontally (e.g. top-right to top-left) or vertically (e.g., top-right to bottom-right) from its starting position. Two directions of travel crossed with four starting positions yielded eight conditions in total (see Figure 3.1D). As before, the probe appeared either when the player was crossing left-to-right (the seventh crossing) or right-to-left (the eighth crossing).

In the post-game survey, subjects were asked about the motion of the unexpected object and its appearance, but were not asked where on screen the object appeared.

Results and discussion

We excluded data from 251 subjects (23% of our sample) from our analysis using the same criteria as prior experiments.

In this experiment, to be counted as having noticed the unexpected object for the primary analysis, subjects had to (a) report having noticed a new object, (b) report that it was moving, and (c) correctly identify its direction of motion from a choice of five directions (up, down, left, right, or not moving).

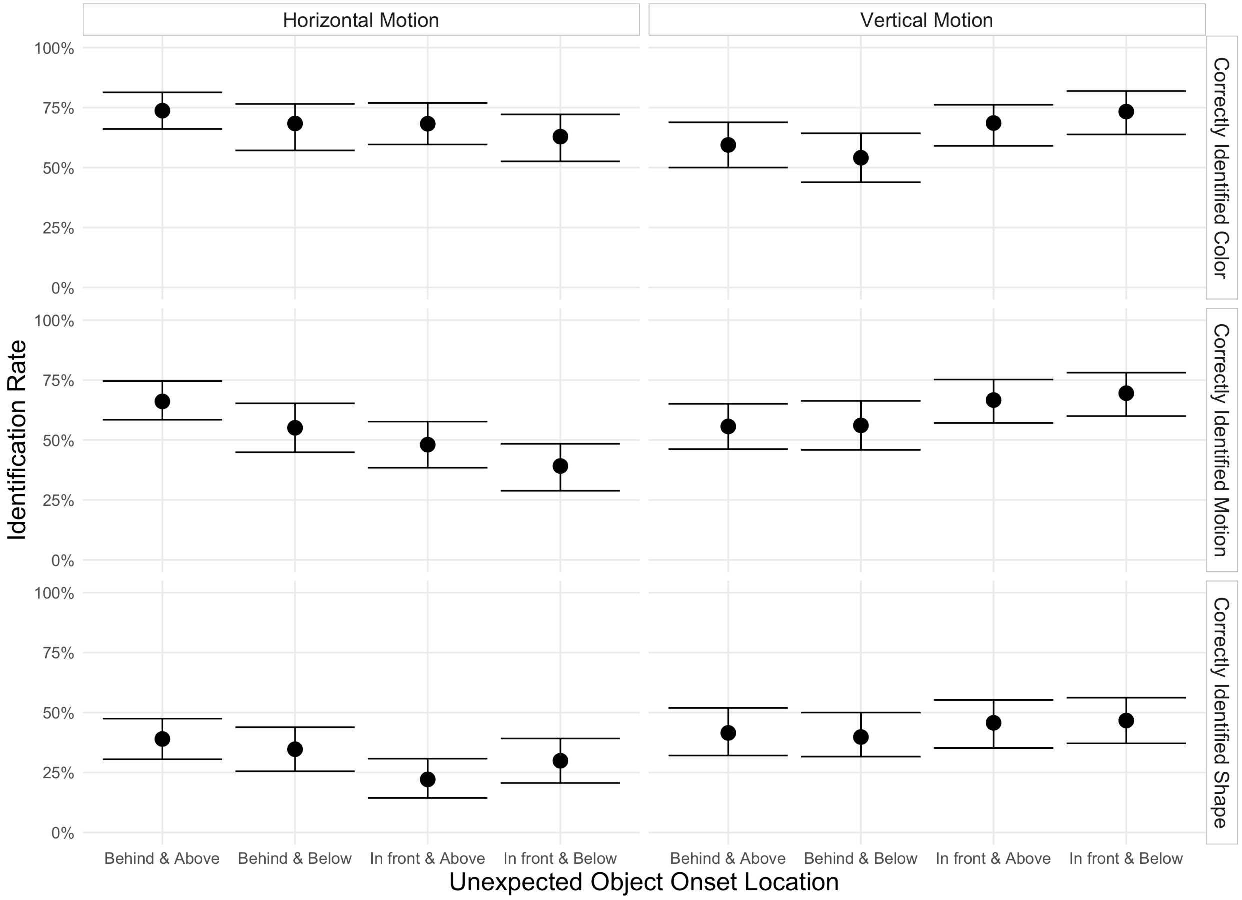

Figure 3.4. Rate at which subjects who reported seeing a new object correctly identified the unexpected object's features, broken down by each possible position and motion trajectory. Error bars are 95% bootstrapped confidence intervals. To be counted as correctly identifying a feature of the unexpected object, subjects had to report noticing something new, and: for color, report that the new object was green; for motion, select the correct trajectory; and for shape, report that it was a diamond.

Collapsing across motion direction, we observed similar location effects as Experiment 2. Unexpected objects that traveled horizontally above the subjects' avatar were noticed 10.5 percentage points (95% CI: [0.8, 20.4]) more than the objects that traveled horizontally below the avatar. Objects that traveled vertically in front of the avatar had a 12.2 percentage point advantage (95% CI: [2.5, 21.8]) over objects that traveled vertically behind the avatar. These results replicate the patterns observed in Experiments 1 and 2 with the static object locations, once again indicating that attention is allocated according to the constraints imposed by the direction of travel and obstacle avoidance (Figure 3.4).

There was no substantial difference in noticing for objects that traveled upwards from below versus objects that traveled downwards from above when collapsing across position (an overall difference of 1.9 percentage points; 95% CI: [-7.2,11.3]). Although vertically-moving objects that appeared in front of the avatar were noticed more than those that appeared behind, the upward and downward trajectories were noticed at similar rates in each case (a 2.9 percentage point difference between upwards and downwards trajectories for the in-front objects, 95% CI: [-9.5, 14.3], and a 0.5 percentage point difference for the behind objects, 95% CI: [-13.4, 13.3]).

There was a difference in noticing for objects that started behind the avatar and overtook it as they traveled horizontally compared to those that started in front and traveled towards the avatar (an overall difference of 17.3 percentage points; 95% CI: [7.6, 27.2]). As for the vertical trajectories, this pattern was consistent regardless of position (an 18 percentage point difference in noticing between overtaking and passing objects moving above the avatar, 95% CI: [5.4, 31.1], and a 15.9 percentage point difference for objects below the avatar, 95% CI: [1.6, 29.3]).

Results for vertically moving unexpected objects did not support a difference in noticing when an object moves from an attentionally relevant area into an irrelevant one, or when it moves from an irrelevant region to a relevant one; the only major difference was the overall effect of in-front versus behind that we observed in earlier experiments.

For horizontally moving objects, the results appear consistent with greater noticing of objects that move into a relevant region from an irrelevant one, given that noticing rates were higher when the unexpected object started behind and traveled alongside the avatar. However, that pattern of motion also meant that the unexpected object spent more time near the player's avatar if the avatar moved while the unexpected object was onscreen. While the time in front versus behind the avatar was equated, the objects that traveled towards the avatar spent much less time nearby than the one that tracked alongside it and overtook it. The large difference in noticing could be due entirely to this difference in proximity. Although motion direction is confounded with proximity within a position, we nevertheless observed the same overall above versus below advantage that we saw in previous experiments when collapsing across these motion directions.

Experiment 4 attempts to replicate the critical finding of greater noticing when an object moves from an irrelevant to a relevant region while controlling for the confound of time nearby the player's avatar.

EXPERIMENT 4

In Experiment 4, we used unexpected objects whose trajectories and distance to the subjects' avatar were equated across conditions, varying only whether an object started outside of the assumedly attended region and moved into it or vice-versa. Finding a large difference in noticing of the unexpected object between the conditions (as in Experiment 3, but without the proximity confound) would indicate an effect of the unexpected object's trajectory into or out of an attentionally relevant area on noticing.

Methods

A demonstration of the task may be viewed at http://simonslab.com/game/xtransit_demo.html.

Subjects. We anticipated a 20% exclusion rate, so we recruited 291 subjects to finish with 100 per condition.

Materials and procedure. Methods and gameplay were identical to those described in the General Methods, except for a change in the motion of the unexpected object. The unexpected object appeared when the subject has traveled 360 pixels, and always appeared 122 pixels behind the player horizontally. It could start in one of two vertical locations; 122 pixels above the player, or 244 pixels above the player (Figure 3.1E). The object appeared when the player was either crossing left-to-right (crossing 7 of 10) or right-to-left (crossing 8 of 10).

After onset, the unexpected object moved diagonally, traveling at 4 pixels per second in the x-dimension and 2 pixels per second in the y-dimension, traveling 244 pixels horizontally and 122 pixels vertically total. If the unexpected objects started “far” above the player's avatar (244 pixels), it moved diagonally downward to overtake the player and finish close to them (122 pixels above and 122 pixels in front). If it started near to them (122 pixels above), it moved diagonally upward to finish farther away from them (244 pixels above and 122 in front). The two possible motion paths are reflections of each other, so distance to the player over the course of the trajectory was identical (assuming that the subjects either (a) moved at a constant rate while the probe was onscreen or (b) or that players in the two conditions had similar patterns of motion while the probe was on screen). This manipulation therefore controlled for the amount of time spent nearby the player’s avatar while allowing us to test whether an unexpected object that moves into a more relevant area (the area above and in front of a player) is noticed more often than one that moves into a less relevant area (farther above the player).

Due to the unexpected object’s diagonal trajectory, when subjects were asked to report the object’s motion, they were required to select the direction they thought it moved from eight arrows (four pointing to the cardinal directions, four to the inter-cardinal directions).

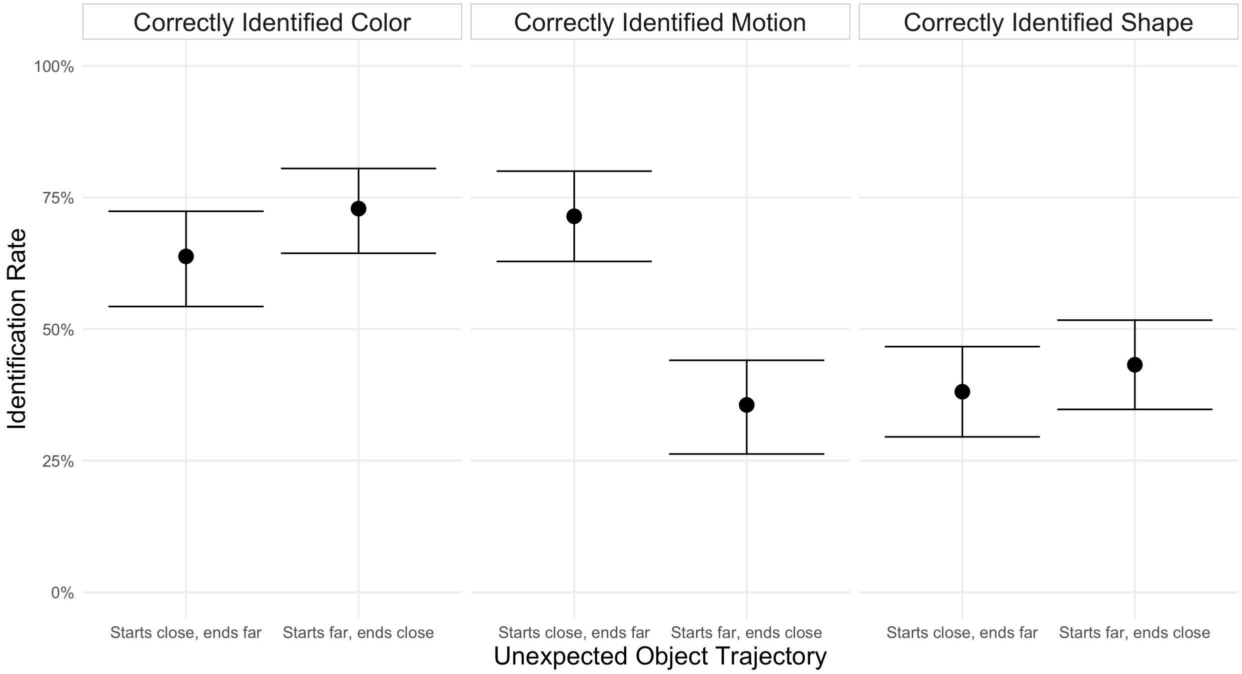

Figure 3.5. Rate at which subjects who reported seeing a new object identified the unexpected object's features, broken down by the unexpected object's trajectory. Error bars are 95% bootstrapped confidence intervals. To be counted as correctly identifying a feature of the unexpected object, subjects had to report noticing something new, and: for color, report that the new object was green; for motion, select the correct trajectory from the cardinal and inter-cardinal directions; and for shape, report that it was a diamond.

Results and discussion

We excluded data from 68 subjects from analysis (23% of our sample). As in Experiment 3, our primary criterion for noticing was correct identification of the unexpected object's motion. Subjects had to report noticing something new, report that it was moving, and choose the correct direction of motion from an array of arrows.

There was a large difference in noticing between conditions. Subjects noticed an unexpected object that appeared near them and moved away 71.4% (95% CI: [62.9, 80.0]) of the time, but noticed an object that appeared far from them and got closer only 35.6% (95% CI: [27.1, 44.1]) of the time (a difference of 35.8 percentage points, 95% CI: [23.8, 48.2]; Figure 3.5). Noticing rates were similar to the approximately comparable condition in Experiment 3, in which the unexpected object started above and behind the avatar and traveled horizontally to overtake it (noticed 66% of the time).

Unexpectedly, and unlike in previous experiments, the pattern of correct identification between conditions varied across features. Although the starts-close, ends-far group was nearly twice as accurate at identifying the unexpected object's motion as the starts-far, ends-close group, the size of the difference was not just smaller for identification of color and shape, but in the opposite direction. The starts-far, ends-close subjects correctly identified the unexpected object's color 9.1 percentage points more than the starts-close-ends-far group, (95% CI: [-2.9, 21.2]) and correctly identified the shape 5.1 percentage points more (95% CI: [-7.5, 18.4]). Why do these groups differ in their ability to identify the motion direction, but less so in their ability to identify other features of the the unexpected object?

One possibility is that the time course of noticing differs between the two conditions. Subjects may notice the unexpected object once it draws near. If so, when it starts nearby and travels away, subjects would notice it sooner and be able track it during the entire course of its movement. In contrast, when the object starts far away and gets closer, they may not notice it until the last moment and cannot track its path of motion over time, but can identify its other features.

EXPERIMENT 5

Experiment 5 tested whether the timing of noticing might explain the difference in motion identification between the two conditions. The study duplicated Experiment 4 with a change to allow us to determine roughly when subjects noticed the unexpected object: the unexpected object changed color halfway through its trajectory. If the difference in accuracy when reporting the unexpected object's motion between the two conditions in Experiment 4 was due to noticing the object early versus late, we should see more subjects reporting the unexpected object's second color in the condition in which the unexpected object appears far away and gets closer.

We expect to observe the same pattern of unexpected object feature identification as in Experiment 4; a large difference between conditions in correct identification of the motion of the unexpected object, but no such differences for shape or color identification. If the difference for motion identification results from when the unexpected object is noticed, then we should find that the subjects who reported noticing the unexpected object and could correctly identify its color are more likely to report the earlier color when the unexpected object onsets close to the avatar, and the later color when the unexpected object onsets far away from the avatar.

Methods

A demonstration of the task may be viewed at simonslab.com/game/xcol_demo.html.

Subjects. We recruited 313 with the goal of 100 usable subjects per condition.

Materials and procedure. The gameplay was identical to that described in the General Methods; however, the number of required crossings was reduced from 10 to 8. The median time to complete the game in Experiment 4 was roughly 4.5 minutes. In order to maintain a fair pay rate for the task, the gameplay portion was shortened. The unexpected object thus appeared randomly on the 5th or 6th of 8 crossings; the procedure was otherwise unchanged.

The behavior and movement of the unexpected object was identical to Experiment 4. However, the unexpected object started with one of two colors, green (#1bad1b) or yellow (#cccc26). It remained that color for 24 frames, then linearly interpolated to the other color (yellow if it began as green, green if it began as yellow) over the course of 10 frames, then remained its final color for 24 frames before offsetting (Figure 3.1F). In the post-game survey, rather than being asked what color the unexpected object was, subjects were asked what color it was when they first noticed it. All other questions were unchanged from Experiment 4.

Results and discussion

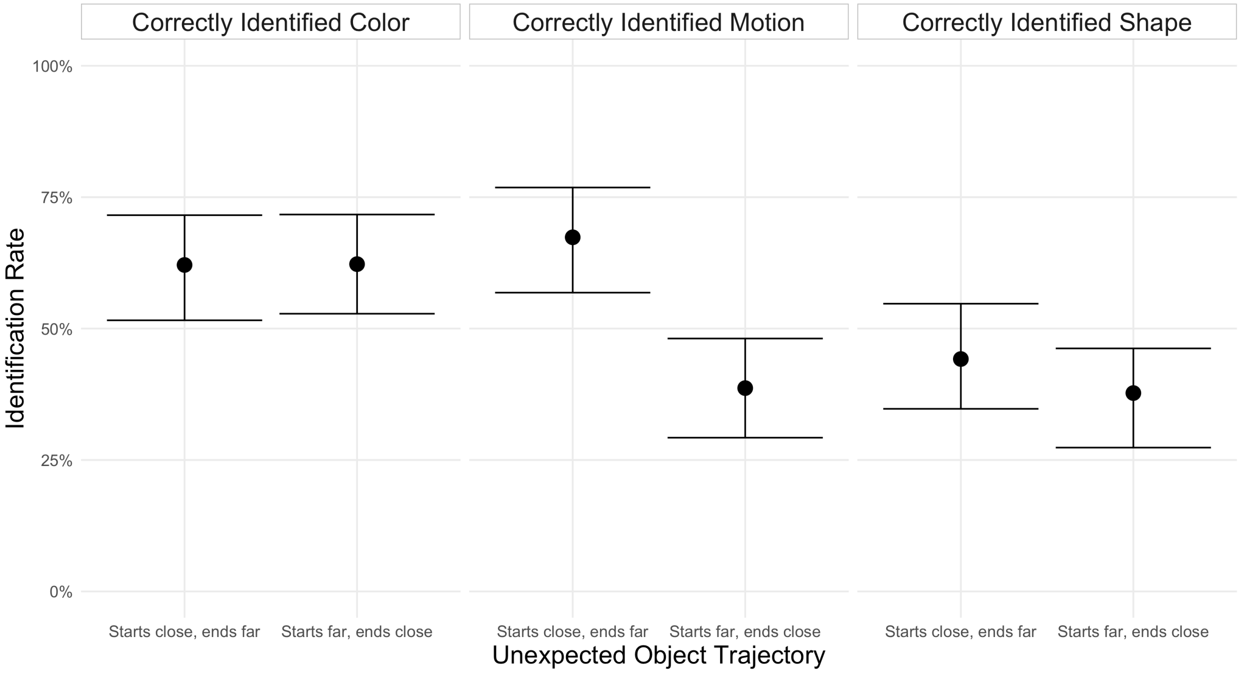

Figure 3.6. Rate at which subjects who reported seeing a new object identified the unexpected object's features, broken down by the unexpected object's trajectory. Error bars are 95% bootstrapped confidence intervals. To be counted as correctly identifying a feature of the unexpected object, subjects had to report noticing something new, and: for color, report that the new object was green or yellow; for motion, select the correct trajectory from the cardinal and inter-cardinal directions; and for shape, report that it was a diamond.

Prior to analysis, we excluded data from 112 subjects (36% of our sample). Overall, we replicated the results of Experiment 4. Correct identification of the path of motion differed starkly between conditions, at 67.4% when it appeared nearby and moved away versus 38.7% when it appeared far away and approached (a 28.7 percentage point difference, 95% CI: [15.8, 41.1]), but correct identification of the other features did not differ much between starts-close, ends-far and starts-far, ends-close (a difference of -0.2 percentage points, 95% CI: [-13.2, 13.2], for color and 6.5 percentage points, 95% CI: [-5.9, 19.8], for shape; Figure 3.6).

Experiment	**Motion Type**	**Identified Motion**	**Identified Vertical Motion**	**Identified Horizontal Motion**
Experiment 4	Starts far, ends close	35.6%	39.0%	73.7%
Experiment 4	Starts close, ends far	71.4%	73.3%	73.3%
Experiment 5	Starts far, ends close	38.7%	41.5%	63.2%
Experiment 5	Starts close, ends far	67.4%	70.5%	70.5%

Table 3.3. Identification rates for the component motion by condition for Experiments 4 and 5. To be counted as noticing the motion overall, subjects had to get the motion direction correct. For the vertical component, they simply had to supply any direction that contained the correct vertical direction (e.g. ‘up-left,’ ‘up,’ or ‘up-right’ would be accepted), and for the horizontal component, any direction that contained the correct horizontal direction.

Among subjects who correctly reported the color of the unexpected object, there was no difference between conditions in the likelihood of reporting the first versus last color. For the start-far, end-close condition, of those who correctly identified one of the object's colors, 33.3% reported the first color and 66.7% reported its second color. For the start-close, end-far condition, 32.2% reported the first color and 67.8% reported the second color. While the color subjects reported is not a perfect indicator of when they noticed the object—the second color may overwrite the first in memory, for instance, or some subjects who see both colors may be biased to report the last color they saw—the absence of a difference between the two conditions likely rules out large differences in the time course of noticing as the explanation for the discrepancy in motion identification.

Why, then, is one group much less accurate than the other in identifying the motion of the unexpected object if there is no difference in when they first notice it? One possibility is suggested by examining the vertical and horizontal components of the motion identification separately (Table 3.3). While subjects in the starts-close, ends-far condition were equally accurate on identifying the horizontal and vertical components of the unexpected object's motion (that is, they reported that it was moving upwards just as accurately as they reported it moving right or left), subjects in the starts-far, ends-close condition were nearly twice as accurate at reporting the horizontal component compared to the vertical one. A post-hoc analysis of the motion reporting in Experiment 4 reveals an identical pattern; the reason the starts-far, ends-close group had such low accuracy is that they were much less likely to detect the vertical component of the motion. Because the unexpected object covers twice as much distance horizontally as it does vertically, the signal may be stronger for the horizontal component of the motion. The starts-far, ends-close group may merely be more uncertain, and guesses the horizontal direction (of which they might be more sure) and disregards the vertical component of the motion. This difference in groups may therefore simply reflect different response strategies under different levels of certainty, rather than any differences in attention.

GENERAL DISCUSSION

The spatial allocation of attention conforms to the demands of the environment, even when that environment is a simple road-crossing game. When the direction of travel is restricted, and people can only travel forward, they are more likely to notice unexpected objects that appear in front of them than behind them, and are more likely to notice nearby unexpected objects than faraway ones. The hazardous objects in the game also play a role in directing attention; subjects were most likely to notice an unexpected object that appeared in front and above them—the area of the display in which the hazards posed the greatest threat. Unexpected objects were less likely to be noticed if they appeared the same distance away from the subjects' avatar but were underneath it, corresponding to the area of the display in which the hazards could no longer collide with the subject's avatar.

The way subjects allocated attention in this task reflected their appraisal of the their ongoing actions and the display environment and not a strategy of searching for the unexpected object. Subjects were not told where to direct attention, were not informed whether they should attend to or ignore any objects in particular, and were not informed about the possibility of additional objects in the display. Even though subjects were free to approach the task however they liked, attention was concentrated to the most task-relevant areas. Not only does this show the role of environmental constraints on attentional allocation, but also demonstrates a naturalistic way to control the spatial deployment of attention without explicit direction.

While the results from the static objects reveal a clear pattern in the spatial allocation of attention in response to the environment, the data from the moving objects indicate little, if any, role of movement through these areas on noticing. Experiments 3-5 attempted to investigate the impact on noticing of objects traveling into and out of attentionally relevant areas. In Experiment 3, unexpected objects moving on vertical trajectories were noticed at the same rate, regardless of whether they onset in the attentionally relevant areas above subjects' avatar and traveled to the less relevant area below the avatar, or vice-versa. Although we observed a difference in noticing for the horizontal trajectories, these were confounded with proximity to subjects' avatar. This difference disappeared after controlling for this confound (Experiments 4-5). When proximity to the avatar was equated over the unexpected object's trajectory, subjects were equally likely to notice it whether it started in a less relevant area of the display and finished in a more relevant one or vice-versa. Although there was a difference in subjects' ability to correctly identify the object's direction of motion between conditions, this appeared to be related to response strategies under uncertainty rather than any meaningful difference in attention. Overall we replicated the findings for static unexpected objects with moving objects, finding more noticing for objects in front than behind, and for above than below. Across experiments, the movement behavior of the unexpected object in the display had much less of an impact on noticing than did the general region of the display in which it appeared. It does not seem to matter whether an object moves into or out of an attentionally relevant area, and something unexpected entering a closely monitored area does not attract any more attention than something leaving it.

Overall, the environment and the demands of performing the road-crossing task shaped the allocation of attention. People tend to monitor the highest-risk areas the most, and pay less attention to areas they cannot access and areas that no longer pose a threat to their actions. While unexpected objects that share features with threatening objects do not seem to be noticed more often than objects sharing features with neutral or rewarding objects in a game context (Stothart, et al., 2017), threatening objects do seem to influence the spatial allocation of attention. Future studies can examine the relative contributions of object features and object locations to noticing in this sort of interactive environment. The nature of the games themselves may determine these contributions; both of these games had a strong spatial and hazard-avoidance component, which might have led participants to prioritize attending to object locations over object features.

Attention operates in a context. Most of the time, we deploy selective attention in the service of a goal. We might expect that the interaction of the structure of the environment, how we navigate through it, and what we intend to accomplish influences both how we deploy attention and what information we select versus filter from awareness. If we want to understand attention in natural tasks like driving or walking, an important first step is exploring attention in smaller-scale, easy-to-control environments. In order to draw conclusions that might generalize to more complex settings, however, we should try to avoid adding constraints that might alter how attention is deployed. For example, dual-task designs in which subjects navigate an environment while also responding to some secondary, unrelated task might mis-measure how we direct attention in the absence of such secondary goals. Inattentional blindness paradigms measure attention while subjects engage more naturally with a display or task without adding extraneous demands on attention, while still providing a naturalistic measure of what people notice.

CONSTRAINTS ON GENERALITY

Space-based effects similar to those we investigate here have emerged in other inattentional blindness paradigms, both run in in person (Most et al., 2000) and on Mechanical Turk (Stothart et al., 2015). We expect the overall effects we found to generalize to any task with similar constraints, and to generalize to in-person, lab-based, or online testing settings, although the particulars of the effects—what areas are emphasized, what distances are monitored, and overall noticing rates—likely will vary according to how the game environment is set up and the particulars of the navigation constraints and obstacle or hazard avoidance.

CHAPTER 4: TEMPORAL DYNAMICS OF NOTICING

This chapter describes previously published experiments.[4]

The work described in Chapters 2 and 3 indicate than when selective attention is narrowly engaged, suppression acts broadly across space and object features to exclude anything other than the attended object. People often find sustained inattentional blindness so remarkable because salient, obvious objects fail to break through, but in light of these findings, ought we be more surprised that objects get through at all? Why does suppression sometimes fail to filter task-irrelevant objects?

A clue to this question may lie in *when* unexpected objects reach awareness, should they manage to do so. There are various possibilities, each of which suggests a different mechanism by which suppression is operating in these sustained inattentional blindness paradigms. If noticing is tethered exclusively to onset or offset, for instance, it suggests that something about the visual change to the display provides an opportunity for a new object to slip past the filters. If objects break through completely at random, it suggests that suppression is not a complete blackout, but a stochastic back-and-forth between the distracting information and the attempts to filter it. At any moment, the suppression could wane and allow the unexpected object through. Unexpected objects could be more likely to break through the longer they are visible, indicating some sort of information accumulation that eventually reaches awareness.

Although it is a simple question on its face, knowing when an unexpected object reaches awareness can offer profound insights into the nature of the suppression in inattentional

---

[4] Wood, K., & Simons, D. J. (2019). Now or never: Noticing occurs early in sustained inattentional blindness. *Royal Society Open Science, 6*(11), 191333. doi: 10.1098/rsos.191333. Published and reproduced under a CC-BY license (http://creativecommons.org/licenses/by/4.0/).

blindness. In the following experiments, subjects were presented with unexpected objects for varying exposure durations. Subjects also reported where in the display they first detected the object. Overall noticing rates increased only slightly, even with large increases in exposure time, and subjects who noticed the object consistently localized it near onset. This suggests that unexpected objects have a brief window of time after onset to reach awareness, after which they are suppressed and unlikely to draw notice even with extended time in the display.

INTRODUCTION

As anyone who has ever tried to get the attention of a distracted friend knows, people can be remarkably oblivious to new or unexpected events when they are sufficiently engrossed in something else. Someone can be inattentionally blind to your waving (and inattentionally deaf to your calling out). A large body of work has attempted to delimit when people will experience inattentional blindness (Mack & Rock, 1998), what aspects of the unexpected event affect noticing, and how engagement in a primary task matters. Fewer studies have examined an equally interesting question: When do people notice? Is your friend guaranteed to notice you eventually if you keep waving? When during your efforts will they be most likely to spot you?

Most studies of the temporal characteristics of selective attention have focused on tasks in which subjects know they will have to attend to some areas or objects and ignore others; they can evaluate and establish attentional filters for all relevant aspects of the display and use those filters repeatedly across many trials. In these circumstances, attention can be deployed to a particular location or feature in a few hundred milliseconds (Liu, Stevens, & Carrasco, 2007). For example, a stimulus in a search or RSVP task can be processed in something on the order of

50 ms (Egeth & Yantis, 1997) and inhibition of distractor features can begin as early as 100ms after stimulus onset (Moher, Lakshmaman, Egeth, & Ewan, 2014).

These and other methods of studying the temporal characteristics of attention (e.g., ERPs) cannot be employed to study noticing in inattentional blindness due to the one-shot nature of the phenomenon. Once subjects become aware that something unexpected may appear, it is no longer unexpected.[5] Because it is essential that subjects cannot prepare for the unexpected stimulus, inattentional blindness tasks typically use one trial per subject which precludes the within-subjects approaches of other tasks.

Across different types of tasks, subjects can miss unexpected objects across a variety of time scales. A substantial proportion of subjects—25% to 75%, depending on the precise nature of the experiment—miss a small square flashed for 200 ms in a static display (Mack & Rock, 1998). About half of subjects also miss a cross drifting across the screen for 5 seconds in a multiple object tracking task (Most, Simons, Scholl, Jiminez, Clifford, & Chabris, 2001) or a gorilla striding through a basketball game for 9 seconds (Simons & Chabris, 1999). However, few studies have systematically examined the effect of exposure time on noticing within a single task. When people do notice an unexpected object, at what point during the task does it reach awareness? Would longer exposure to the same object lead to more noticing?

One of the original studies of selective looking, in which participants viewed the same display for different lengths of time, offers a potential answer to these questions (Becklen & Cervone, 1983). Subjects viewed a video of two basketball teams, with instructions to track the passes made by one team while ignoring those made by the other team. During the action, a

---

[5] While repeated inattentional blindness in a single task can occur, the proportion of non-noticing subjects drops substantially with each repeated presentation of the "unexpected" object (Ward & Scholl, 2015).

woman holding an umbrella walked through the display. In one condition, the woman passed all the way across the screen and was visible for about 5.5 seconds. In another condition, the experimenters stopped the video when she was about halfway across, amounting to roughly 2 seconds of exposure. The primary purpose of the study was not to examine the effect of exposure on noticing, but it did appear to make a difference: 34% of subjects in the 5.5s condition noticed the umbrella woman, whereas just 7% did in the 2s condition. More exposure to the unexpected object apparently increased the probability that it would be noticed. As the study itself noted, however, there may have been critical content differences between the first 2 seconds of the video and the remaining 3.5 seconds (other than just exposure time) that contributed to differences in noticing.

More recent studies examined the influence of exposure time on noticing with a multiple object tracking task in which subjects monitor a subset of objects and count the number of times they bounce off the edge of the display (Beanland & Pammer, 2010a). One experiment found little difference in noticing for a "fast" unexpected object that took 5 seconds to cross the display and a "slow" one that crossed in 9 seconds. However, the study used relatively small samples (*n* = 25 per condition) with low rates of noticing for the first appearance of the unexpected object (4 and 5 noticers in the slow and fast conditions respectively), so the study does not permit definitive conclusions about whether or not exposure time mattered for noticing.

Another recent study also varied exposure time to the unexpected object by varying its speed in a multiple object tracking task (Kreitz, Furley, & Memmert, 2016) and observed higher noticing rates for the slower object that was on screen for longer. Noticing rates were comparable when the fast and slow objects were onscreen for the same amount of time, suggesting that increased exposure to the objects, rather than the speed difference), increased the likelihood of noticing.

A much earlier and more unusual study, intended as a test of how people perceive and respond to seemingly paranormal events, provided early evidence of inattentional blindness for real world events and also collected data on approximately when observers noticed the unexpected event (Cornell, 1959). The experimenter dressed in a sheet and walked back and forth across the stage of a movie theater while a trailer was playing before the film. The “ghost” was visible for 50 seconds, and 32% of the theater audience did not report seeing it. Of the 68% who did notice the ghost, just over half saw it in the first 5 seconds it was visible (inferred from the part of the ghost’s walk that they reported).

The methods used in each of these studies provide different information about the time course of noticing. Varying the amount of time the unexpected object is visible tests how much noticing increases with additional exposure time. Asking when subjects first noticed the object narrows down when noticing occurs, particularly if noticing does not vary between exposure times. If the same proportion of subjects notice the unexpected object with a long and short exposure, the location reports can clarify why (e.g. noticing occurs at offset). In order to get a more complete picture of the time course of noticing, we can collect both kinds of data in the same experiment by varying the amount of time the unexpected object is onscreen, keeping all other aspects of the display identical, and asking subjects to report where the object was when they first noticed it. Subjects are fairly accurate at localizing an unexpected object when they do notice it (Newby & Rock, 2001), and the location reports should be sufficient to disambiguate cases where the noticing rate does not vary with exposure time.

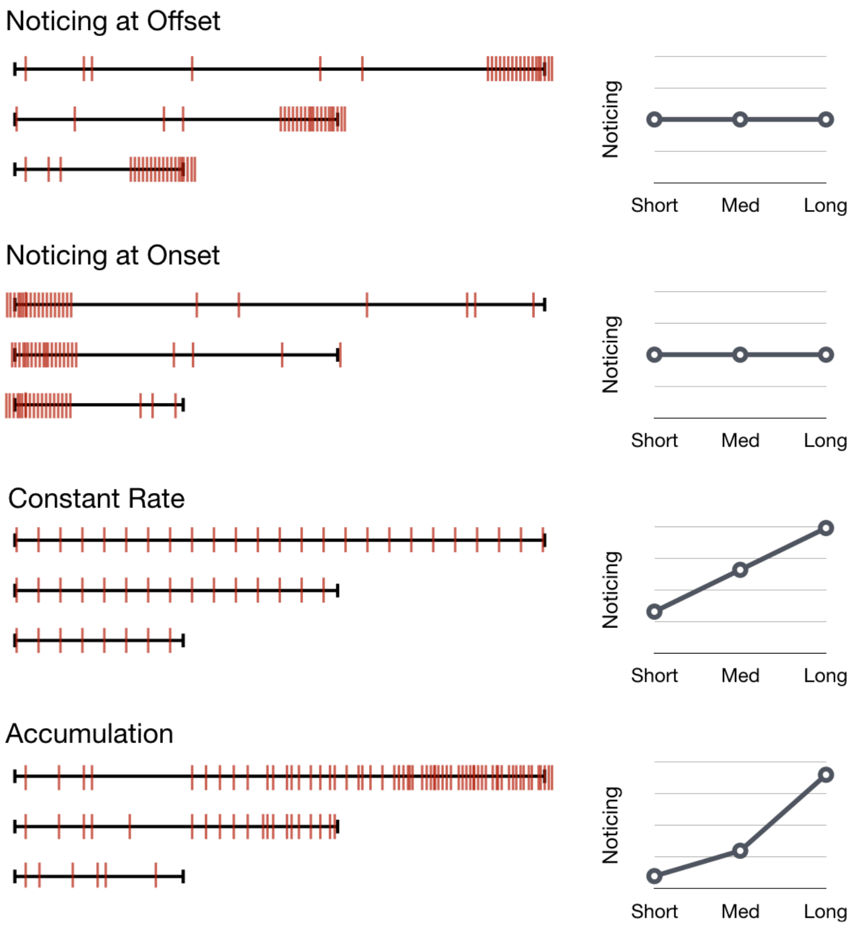

Figure 4.1. Example models for noticing events. The timelines on the right are sample distributions of location reports, with onset on the left and offset on the right. The plots on the left show the corresponding shape of the noticing rates across time intervals. Some models, such as noticing at onset/offset, predict a constant rate of noticing across time, but different patterns of localization. Others predict a change in noticing across time intervals.

If noticing is triggered by a transient event such as the onset or offset of the unexpected object, then noticing rates should be constant across exposure times, and the location reports should cluster around the location of that onset or offset (Figure 4.1). If noticing is a stochastic process (see Kreitz, Furley, Memmert, & Simons, 2015), greater exposure time should provide more opportunity to notice, so noticing rates across participants should be higher with increasing exposure time (as observed by Becklen & Cervone, 1983 and Kreitz, Furley, & Memmert, 2016). That pattern would match the intuitive idea that if your friend doesn't see you at first, it will help to keep waving.

GENERAL METHODS

Subjects

The University of Illinois Institutional Review Board (IRB) waived the requirement for signed consent due to the low-risk nature of the experiment. Prior to accepting the HIT ("Human Intelligence Task," the term for the jobs posted to Amazon's Mechanical Turk service), subjects were shown an information screen that provided experimenter and IRB contact information. It explained that their responses would be anonymous, described how their data would be used, and noted that their participation was voluntary. All subjects were US-based workers recruited through Amazon's Mechanical Turk ("MTurk") service who had completed at least 100 HITs and had a HIT approval rating of at least 95%. We used TurkGate (Goldin & Darlow, 2013) to exclude subjects who had previously participated in experiments from our lab based on their worker ID. Subjects were directed to an external website running the experiment in Javascript, and upon finishing the experiment, they received a completion code which they entered on MTurk to receive payment ($0.30).

Subjects were automatically recruited in batches of up to nine using the boto3 Mechanical Turk SDK (https://github.com/boto/boto3). When we passed the recruitment threshold for an experiment, recruitment stopped and no further HITs were posted. Based on previous experiments using a similar procedure and small-sample pilot studies for this project, we anticipated a 50% exclusion rate.

Materials and procedure

All experiment methods, analysis plans, and code were preregistered on the Open Science Framework (OSF; https://osf.io/gb6v5/). Each experiment was preregistered separately, prior to

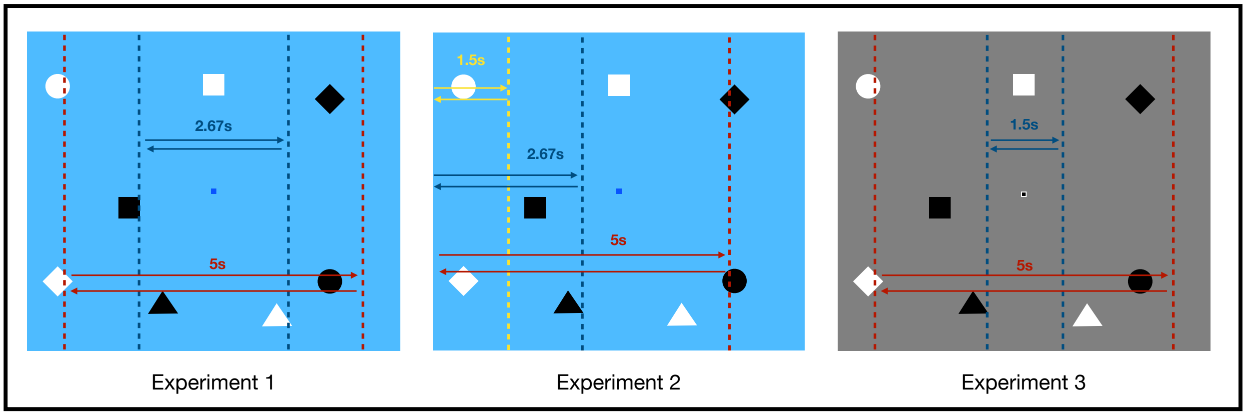

Experiment 1 Experiment 2 Experiment 3

Figure 4.2. The paths taken by the unexpected object in each experiment. The unexpected object always traveled centered in the display. The dotted lines indicate the edges of the invisible occluders that the unexpected object emerged from and disappeared behind. Note that in Experiment 2, the depicted trajectories could also be positioned at the right edge of the display.

data collection for that experiment. Each experiment's project page includes anonymized data, all experimental materials and code, preregistration documentation, and demo versions of the task. Pilot data is also available on OSF, where applicable.

Upon accepting the HIT and navigating to the external website, subjects viewed an instruction screen explaining that they would see a group of eight objects, four white and four black. Subjects were randomly assigned to attend white or black, and were told to count how many times the "attended" color of object bounced off the edges of the display area (ignoring the bounces of the other objects). They were advised to fixate on a central fixation square throughout the task.

After reading the instructions, subjects proceeded to three trials of counting bounces. The display window was 600 x 700 pixels. It was light blue (#58ACFA) in Experiments 1 and 2 and mid-gray (#808080) in Experiment 3 (Figure 4.2). A 10 x 10 pixel fixation square was centered in the display, dark blue (#0000FF) in the first two experiments and black with a white border in the third. The display contained four black objects and four white objects. Each set featured a square (40 by 40 pixels), a triangle (50 pixel base, 50 pixel height), a diamond (56 pixel width by 56 pixel height), and a circle (46 pixel diameter). The objects were randomly placed in the display area at the start of each trial and remained stationary for 1 second before they began moving to give subjects time to prepare for the start of the trial.

The objects moved between 66 and 198 pixels per second and randomly increased or decreased their velocity by 66 pixels per second after a randomly selected period ranging from 300 to 1000 milliseconds, with the constraint that objects never moved faster than 198 pixels per second or slower than 66 pixels per second. Objects occluded one another when they crossed paths, but always passed behind the fixation square. When objects came into contact with the edge of the display, they "bounced" off at a 45 degree angle from the edge at the same speed, with a reversed horizontal direction if they contacted the left or right edge of the display and a reversed vertical direction if they contacted the top or bottom edge. After 15 seconds of motion, the trial ended and subjects were prompted to enter their count of the bounces into a text box which only permitted integer responses.

On the third trial, an additional unexpected object passed through the display. Although some parameters varied by experiment (see each experiment's method for details), the object was always a cross (40 x 40 pixels with arms 14 pixels thick) and traveled at 132 pixels per second, approximately the average velocity of the display objects. It always offset with 2 seconds remaining in the trial, although the amount of time it was onscreen and the location at which it onset and offset varied by experiment. It always traveled horizontally through the display at the vertical midpoint, and whether it crossed left-to-right or right-to-left was random for each subject.

Following the third trial, subjects entered their bounces as usual, but were then asked whether they had noticed anything new on the last trial that had not appeared in earlier trials. Next they were asked to indicate the new object's shape from a drop-down menu of options and its color, either from a drop-down menu (Experiments 1 and 2) or a continuous color slider (Experiment 3). Subjects were then presented with a 2/3rd scale image of the display rectangle and fixation square, as well as a scaled-down version of the unexpected object as it had appeared on the critical trial. The unexpected object's icon started in the upper-left corner of the scaled-down display, and subjects were instructed to move the unexpected object to the point in the display where they first noticed it.

Subjects next reported whether they needed vision correction, defined as "glasses or contacts," and if they were wearing it during the experiment, then indicated any technical issues they experienced. Finally they were asked whether they had any prior experience with inattentional blindness tasks. After completing the questionnaire, they were given the completion code for the experiment. In total, the experiment took most participants approximately 3 - 5 minutes to complete.[6]

Analysis Procedure

Data were analyzed in R (R Core Team, 2019) using the packages dplyr (version 0.7.6; Wickham, François, Henry, & Müller, 2018), purrr (version 0.2.5; Henry & Wickham, 2018), tidyr (version 0.8.1; Wickham & Henry, 2018), ggplot2 (version 3.3.0; Wickham, 2016), viridis

---

[6] We do not typically include full attention trials in online experiments. The original studies of inattentional blindness presented the unexpected object briefly and needed to verify that it was visible when people were looking for it (Mack & Rock, 1998). For these dynamic tracking tasks, there is little concern about visibility; given their size, high contrast, and extended time onscreen, the unexpected objects are well above threshold for visibility. Indeed, other researchers using online experiments to study inattentional blindness do not routinely employ full-attention trials (e.g. Ward & Scholl, 2015; Drew & Stothart, 2016).

(version 0.5.1; Garnier, 2018), ggforce (version 0.2.2; Pedersen, 2019), and circular (version 0.4-93; Agostinelli & Lund, 2017). For all analyses, we report point estimates for values of interest with 95% bootstrapped confidence intervals calculated via the percentile method (Efron & Tibshirani, 1993). For comparisons of interest, we also calculate difference scores and their 95% bootstrapped confidence intervals. All estimates and comparisons were preregistered on OSF.

Exclusion Criteria

Our preregistered criteria excluded data from subjects who reported being younger than 18 years old; whose bounce counts erred by more that 50% in either direction on two or more trials; who reported needing vision correction but not wearing it during the experiment; who reported that the experiment lagged, froze, or had some other technical problem; or who reported prior experience with inattentional blindness tasks. In Experiment 3, we also excluded participants with a confusion index greater than 1.78 on the Farnsworth D-15 task that is designed to measure color vision (Farnsworth, 1947; Vingrys & King-Smith, 1988). For a detailed breakdown of the exclusions in each experiment, see Table 4.1.

Exclusion Rule	**Experiment 1**	**Experiment 2**	**Experiment 3**
Miscounted by more than 50% on more than 2 trials	190	367	333
Reported being younger that 18	0	0	1
Reported needing vision correction but not wearing it	32	105	82
Reported a technical problem with the experiment	107	140	131
Reported prior experience with inattentional blindness	24	37	26
Had a severity index greater than 1.78 on the Farnsworth D-15	NA	NA	182

Table 4.1. The number of subjects excluded under each rule in each experiment. Subjects could have been excluded under multiple rules, so the total for the experiment may not match the raw sum of all exclusions.

Total Excluded	243	517	516
Total Retained	283	756	488
Total Recruited	526	1273	1004

Table 4.1 (cont.). The number of subjects excluded under each rule in each experiment. Subjects could have been excluded under multiple rules, so the total for the experiment may not match the raw sum of all exclusions.

EXPERIMENT 1

Does increasing the exposure time to an unexpected object also increase the likelihood that it is noticed? Although this idea is intuitive and has some tentative support (Becklen & Cervone, 1983), no previous study has systematically compared the effect of exposure times on noticing under otherwise identical conditions. Experiment 1 compared noticing of an unexpected object that was visible for either a long exposure of 5 seconds or a shorter exposure of 2.67 seconds, corresponding to the unexpected object crossing either 80% of the total width of the display or 40% of the total width display. A 5-second exposure is typical for this sort of sustained inattentional blindness task (e.g. Most, Simons, Scholl, Jiminez, Clifford, & Chabris, 2001), and these exposure durations are similar to those used in previous studies of the influence of exposure time on noticing (Becklen & Cervone, 1983). If noticing is a stochastic or accumulative process, with greater time leading to more noticing, then noticing rates should be lower in the 2.67 second condition because those participants have 2.33 fewer seconds to spot the unexpected event; if noticing is instead driven by an onset or offset event, then noticing rates should not differ between the conditions.

In addition to the overall noticing rates, we also assessed where those participants who noticed the unexpected object first saw it. This information provides a more fine-grained estimate of the time course of noticing, as their location reports indicate when they noticed the unexpected object. If most reports fall close to offset, for example, that would indicate a tendency to notice the object late, while clustering near onset would indicate early noticing. In both cases, however, the noticing rates would be the same across exposure durations, since noticing would be triggered by an event common to all presentations (Figure 4.1). To verify that these location reports represent a meaningful signal about when people noticed the unexpected object, we included a condition with no unexpected object in this experiment. This condition will provide the true random baseline against which to compare the localization reports from the other conditions.

Methods

The materials and preregistration for this experiment are available at https://osf.io/yekzc/. A demonstration of the task may be viewed at simonslab.com/mot/temporal_mot_demo.html.

Subjects. We aimed to recruit 500 subjects with the goal of collecting usable data (i.e., after exclusions) from 100 subjects in each unexpected object condition and 50 subjects with no unexpected object. We recruited according to the procedure in the General Methods and ended up with 526 subjects in total.

Materials and procedure. The primary task and post-survey questionnaire were as described in the General Methods.

The unexpected object in this experiment was a mid-gray (#808080) cross. Subjects were randomly assigned to one of three possible unexpected object conditions. In the short-duration condition, the unexpected object appeared onscreen for 2.67 seconds, onsetting 10.33 seconds into the trial by emerging gradually from behind an invisible occluder positioned 210 pixels from the edge of the display and traveling 280 pixels before offsetting behind another invisible occluder. In the long-duration condition, the unexpected object appeared for 5 seconds, onsetting 8 seconds into the trial from an invisible occluder 70 pixels in from the edge and traveling for 560 pixels before offsetting (Figure 4.2). In the no-unexpected-object condition, the procedure was identical to the other two conditions except that no additional object appeared on the critical trial. Subjects had a 2 in 5 chance of being assigned to either unexpected object condition and a 1 in 5 chance of being assigned to the no-unexpected-object condition.

Results and Discussion

Prior to analysis, we excluded 243 subjects (46% of our sample) according to the criteria outlined in the General Methods, leaving 283 in the analysis (*n* = 54 in the no-unexpected-object condition, *n* = 104 in the 2.67s exposure condition, and *n* = 125 in the 5s exposure condition). Data are available at https://osf.io/zmvw2/.

Noticing. As specified in our preregistration, we coded subjects as having noticed the unexpected object if they were assigned to a condition that had an unexpected object, reported noticing something new, and correctly reported either the object's shape or its color.

44.0% (95% CI: [35.2, 52.0]) of subjects in the 5s exposure time condition noticed the unexpected object, and 38.5% (95% CI: [29.8, 47.1]) of subjects in the 2.67s exposure time condition noticed it for a difference of 5.5 percentage points (95% CI: [-8.1, 18.0]; Figure 4.3).

Location reports. For the location analysis, we looked separately at the vertical and horizontal localization of the unexpected object. Recall that the object always traveled along a horizontal path at the vertical midpoint of the display. Although vertical localization of the unexpected object is not informative about when people noticed the unexpected object, it does indicate whether subjects are placing the object in the area it actually appeared. The pattern of

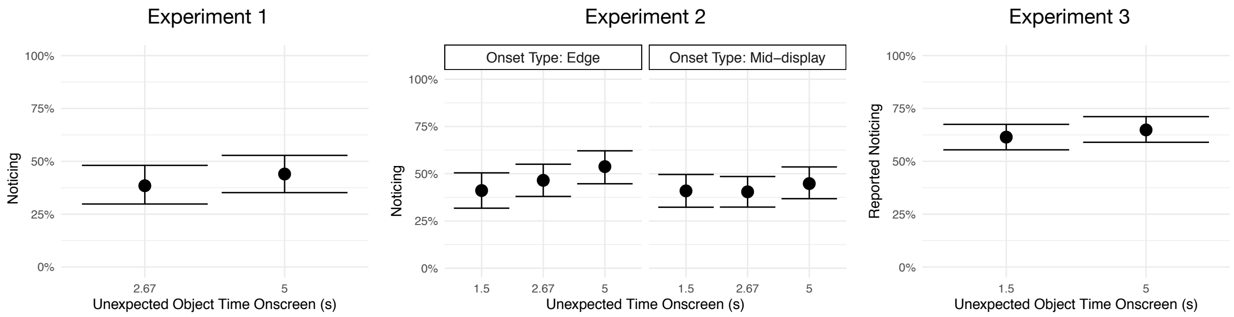

Figure 4.3. Noticing rates or self-reported noticing for subjects in each experiment, broken down by unexpected object exposure time. Error bars are 95% bootstrapped confidence intervals.

localization still differed between noticers and non-noticers. Noticers fairly consistently placed the object near the horizontal midline (*M* = 282 pixels, *SD* = 67; actual midline = 300 pixels).

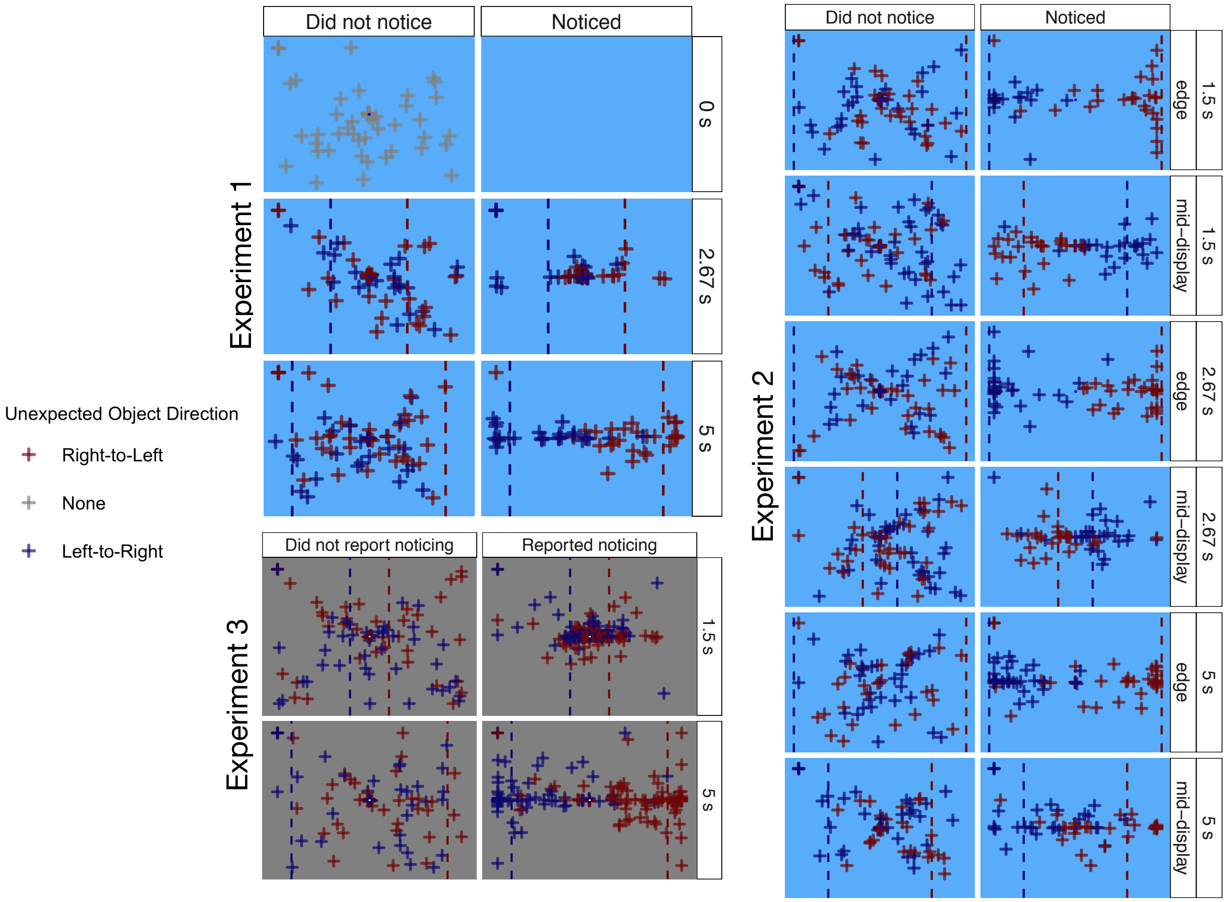

Figure 4.4. Scatterplots of the location reports from each subject in each condition. The dashed vertical lines indicate the onset points in each condition, and are color-coded according to the direction the unexpected object traveled after it appeared. Each panel shows both the left-to-right and right-to-left variant for a particular condition, color-coded accordingly.

The average placement for non-noticers was also near the midline, but with far greater variability (*M* = 313 pixels, *SD* = 120). The average placement for those in the no-unexpected-object condition was also near the midline, but again with greater variability than for notices (*M* = 311, *SD* = 153).

To estimate the timepoint of noticing, the position along the horizontal midline is more critical. Recall that participants were randomly assigned to experience an unexpected object moving from right to left or from left to right. The object placement for noticers was tied to the actual onset location they experienced, with placement clustering near the onset position (Figure 4.4). There is especially clear separation in the 5s condition, when the left- and right-side onset points are further apart (the lower right panel for Experiment 1 in Figure 4.4). 81% of noticers placed the unexpected object on the onset side of fixation. In contrast, non-noticers and subjects in the no-unexpected-object condition tended to place the object around fixation with a large spread (left column for Experiment 1 in Figure 4.4; also see Table 4.2).

Noticed Unexpected Object	Time Onscreen	Unexpected Velocity	Mean Distance of Reports to Fixation (SD)	Mean Distance of Reports to Onset (SD)	Mean Distance of Reports to Offset (SD)
No	0	None	193.2 (138.8)	NA	NA
No	2.67	Right to left	151.5 (133.5)	202.7 (111.6)	228.6 (121.4)
No	2.67	Left to right	127.6 (111.7)	215.4 (115.7)	166.4 (95.9)
No	5	Right to left	165.6 (124.2)	329.3 (156.4)	304.9 (129.8)
No	5	Left to right	158.6 (112.7)	288.7 (125.1)	339.2 (136.9)
Yes	2.67	Right to left	86.3 (85.7)	127.9 (56.8)	194.3 (107.1)
Yes	2.67	Left to right	140.6 (156.1)	155.5 (90.7)	262.8 (159.1)
Yes	5	Right to left	204.4 (116.8)	155.7 (130.3)	460.1 (122.3)
Yes	5	Left to right	184.4 (120.9)	156.8 (115.5)	439.1 (151.3)

Table 4.2. Average Euclidean distance, in pixels, of subjects' location reports for the unexpected object from the onset location, fixation location, and offset location in Experiment 1. For each condition and direction of motion, we averaged the distance between each individual location report and the location of the onset in that condition, fixation, and the offset in that condition.

48% of non-noticers and 33% of subjects in the no-unexpected-object condition placed the object on the onset side of fixation.

Some noticing subjects placed the unexpected object closer to the edge of the display than it ever appeared; these instances are most visible in the 2.67s condition. This may reflect a tendency for some subjects to extrapolate the unexpected object's location based on its velocity; if they notice it coming from the left, for instance, they might over-correct and place it closer toward the left edge than at its actual onset point. This sort of error might reflect an actual misperception of the unexpected object’s location as a result of its motion (Kuhn & Rensink, 2016), noise in the precision of their localization (leading to random spread around the onset point), or a deliberate attempt to put the object where they believe it onset rather than where they actually noticed it. In general, though, subjects who noticed the unexpected object positioned it close to its actual onset location with some precision (Table 4.2).

Collectively, Experiment 1 showed that noticing occurred fairly early after onset, and additional exposure to the unexpected object increased the chance of noticing it only slightly. These results contrast with earlier studies in which reducing the exposure also substantially decreased the proportion of subjects noticing the unexpected stimulus (Becklen & Cervone, 1983; Kreitz et al., 2016).

The idea that noticing occurs soon after onset and does not benefit substantially from additional exposure time is supported by the localization data. The pattern of localization reports for noticing subjects—consistent vertical placement and horizontal placement near the onset location—differed dramatically from that produced by non-noticing subjects and subjects who were not presented with an unexpected object. This difference in localization suggests that these reports are based on detection of the unexpected object and are at least coarsely reliable. Furthermore, there was no obvious difference between the locations reported by subjects who missed the unexpected object and those who were not exposed to one at all, suggesting that non-noticers have not represented anything about the location of the unexpected object. Consequently, we can use their location reports as a random-responding baseline.

Although nearly everything about the task and displays was identical across the two exposure duration conditions, they did differ in one potentially important way: The unexpected object's movement in the 5 second exposure condition spanned most of the display, near to both edges, whereas the motion in the 2.67 second condition was further from the edges. If proximity to the edges of the display influences noticing, that might interact with any effects of time on noticing or with location reports.

EXPERIMENT 2

In Experiment 2, we attempted to replicate and extend the findings of Experiment 1. We added an additional, shorter exposure time of 1.5 seconds to the 2.67 and 5 second conditions from Experiment 1. We also shifted the portion of the display the unexpected object traverses, so that it starts or ends at an edge of the display rather than being centered around fixation. The object could onset from either edge of the display and offset in the middle, or onset near the middle and offset at either edge. This allows us to examine whether there is any effect of an edge versus mid-display onset within an exposure duration. It also permits comparison of the pattern of noticing across durations when the objects onset from the same place onscreen versus when these objects onset at different points, revealing whether the pattern we observed in Experiment 1 changes with onset location.

This manipulation also provides a more robust test of the reliability of the location reports for all display durations. In Experiment 1, the left- and right-side onsets for the short display duration were both near fixation and not well-separated in space. In Experiment 2, even the shortest display duration conditions allow a comparison of an object that onsets at the far left or right edge of the display. Additionally, because there are conditions in which the onset position is in the middle of the display, far from the edges, it will be more apparent if subjects misreport the location of the object. For example, if they tend to extrapolate the location to the edge closest to where the object started (e.g., the left edge if the object traveled from left to right) irrespective of the actual onset location, then location reports should vary only with motion direction. If the misperception is milder, then we might expect a majority of the location reports to overshoot the onset point toward the edge closet to the start of the motion.

Methods

The materials and preregistration for this experiment are available at https://osf.io/jx9vs/. A demonstration of the task is available at simonslab.com/mot/temporal_mot_nc_demo.html.

Subjects. Using the procedures described in the General Methods, we aimed to recruit 1200 subjects in order to end up with approximately 100 in each of six conditions after exclusions. We recruited 1273 in total.

Materials and procedure. The task, questionnaire, and appearance of the unexpected object were identical to those of Experiment 1. Only the behavior of the unexpected object differed.

In Experiment 2, there were six possible conditions representing a full crossing of exposure duration and onset behavior. There were three different exposure times: 1.5 seconds, in which the unexpected object traveled 140 pixels; 2.67 seconds, in which it traveled 280 pixels; and 5 seconds, in which it traveled 560 pixels. There were two different onset behaviors. In the "edge onset" conditions, the unexpected object emerged from one edge of the display and offset behind an invisible occluder positioned 140, 280, or 560 pixels from the other edge edge. In the "mid-display onset conditions, the objects onset from behind an invisible occluder positioned 140, 280, or 560 pixels into the display and offset at the far edge (Figure 4.2); this mid-display-onset condition is similar to Experiment 1, in which objects also onset from behind invisible occluders positioned in the display and away from the edges. Whether the object traveled left-to-right or right-to-left was random for each subject.

Results and discussion

We excluded 517 subjects according to the criteria in the General Methods (41% of our recruited subjects) and retained 756 in the analysis (*n* = 107 in the 1.5s, edge-onset condition; *n* = 127 in the 1.5s, mid-display-onset condition; *n* = 129 in the 2.67s, edge-onset condition; *n* = 136 in the 2.67s, mid-display-onset condition; *n* = 132 in the 5s, edge-onset condition; and *n* = 125 in the 5s, mid-display-onset condition). Data are available at https://osf.io/6fu5v/.

Noticing. Subjects were coded as having noticed the object according to the same criteria used in Experiment 1.

Similar to the pattern observed in Experiment 1, reducing the exposure time had a small effect on noticing (Figure 4.3): 41% of subjects (95% CI: [34.6, 47.0]) noticed the unexpected object in the 1.5s exposure condition, 43.4% (95% CI: [37.0, 49.8]) noticed it in the 2.67 condition, and 49.4% (95% CI: [43.6, 55.6]) noticed it in the 5s condition.

There was also a small difference in noticing between objects that onset near the edge of the display and those that onset near the middle of the display. Edge-onset unexpected objects were noticed 47.6% of the time (95% CI: [42.7, 52.4]) and mid-display-onset objects were noticed 42.0% of the time (95% CI: [37.1, 46.9]). The size of this difference increased with longer exposure times, with a 0.2 (95% CI: [-11.6, 13.3]) percentage point difference between edge and mid-display onsets in the 1.5s condition, a 6.1 (95% CI: [-6.0, 18.0]) percentage point difference in the 2.67s condition, and a 9.0 (95% CI: [-3.1, 22.0]) percentage point difference in the 5s condition.

Location reports. As in Experiment 1, both noticers and non-noticers centered their vertical placements near the horizontal midline but with far more variability among the non-noticers (noticers *M* = 294 pixels, *SD* = 78; non-noticers *M* = 312 pixels, *SD* = 131). Noticers also placed the object on the onset side of fixation far more often than the offset side (85% of participants), whereas non-noticers were more evenly split (53% placed it on the onset side of fixation). Noticers generally placed the object close to the actual onset position, while non-noticer placements were more evenly distributed (Table 4.3). This difference is particularly visible when examining the placements for objects that onset near the left or right edge of the display: Noticers placed the object near the edge of onset, but non-noticers did not (Figure 4.4).

Figure 4 also shows that subjects did not dramatically misperceive the location of the unexpected object. Location reports clustered around onset, even when the onset point was positioned mid-display. Although there was variability in the reports, subjects did not extrapolate the location to the edge closest to where the motion started. Instead, their localizations were varied around the position where it actually first appeared.

The results of Experiment 2 replicate the pattern in Experiment 1. Dramatically reducing exposure time only modestly affected the probability of noticing. Even reducing exposure time by more than two-thirds only reduced noticing by 8 percentage points, suggesting that noticing

Noticed Unexpected Object	Time onscreen	Unexpected Velocity	Onset Type	Distance to Fixation (SD)	Distance to Onset (SD)	Distance to Offset (SD)
No	1.50	Right to left	Edge	176.1 (113.6)	348.8 (152.0)	237.6 (136.3)
No	1.50	Right to left	Mid-display	178.3 (119.5)	240.5 (117.1)	342.4 (143.5)
No	1.50	Left to right	Edge	167.8 (123.3)	365.9 (165.9)	252.8 (144.3)
No	1.50	Left to right	Mid-display	208.6 (125.6)	246.6 (138.2)	341.9 (160.9)
No	2.67	Right to left	Edge	171.1 (133.2)	350.7 (158.9)	182.4 (119.4)
No	2.67	Right to left	Mid-display	159.3 (114.4)	189.0 (111.3)	406.7 (139.5)
No	2.67	Left to right	Edge	147.2 (106.7)	386.3 (141.7)	172.0 (105.1)
No	2.67	Left to right	Mid-display	184.5 (129.7)	183.8 (110.7)	329.3 (129.7)
No	5.00	Right to left	Edge	202.1 (107.3)	383.8 (167.3)	277.5 (141.0)
No	5.00	Right to left	Mid-display	118.9 (100.7)	198.0 (75.5)	416.2 (112.2)
No	5.00	Left to right	Edge	163.0 (92.4)	385.4 (133.8)	250.3 (107.2)
No	5.00	Left to right	Mid-display	190.0 (132.6)	271.3 (112.2)	400.7 (187.6)
Yes	1.50	Right to left	Edge	267.0 (112.8)	202.8 (171.8)	168.9 (133.7)
Yes	1.50	Right to left	Mid-display	174.2 (110.4)	129.0 (67.7)	210.6 (110.1)
Yes	1.50	Left to right	Edge	270.6 (82.0)	103.7 (97.6)	104.4 (59.9)
Yes	1.50	Left to right	Mid-display	184.8 (105.1)	128.9 (101.4)	213.3 (136.2)
Yes	2.67	Right to left	Edge	253.7 (101.0)	130.3 (95.2)	193.0 (89.1)
Yes	2.67	Right to left	Mid-display	129.2 (86.6)	107.7 (84.3)	284.5 (118.6)
Yes	2.67	Left to right	Edge	282.4 (93.4)	126.0 (124.3)	226.8 (76.7)
Yes	2.67	Left to right	Mid-display	147.5 (101.5)	120.2 (106.4)	282.4 (137.8)
Yes	5.00	Right to left	Edge	252.5 (111.8)	218.2 (221.1)	399.5 (152.0)
Yes	5.00	Right to left	Mid-display	163.7 (123.7)	178.8 (104.5)	459.2 (161.6)
Yes	5.00	Left to right	Edge	249.2 (93.1)	119.7 (97.6)	454.8 (98.7)
Yes	5.00	Left to right	Mid-display	200.0 (121.4)	181.2 (112.8)	486.3 (172.8)

Table 4.3. Average Euclidean distance, in pixels, of subjects' location reports for the unexpected object to onset, fixation, and offset in Experiment 2. For each condition, onset type, and direction of motion, we averaged the distance between each individual location report and the location of onset in that condition, fixation, and the location for offset in that condition.

occurs soon after onset if it is to occur at all; there is little additional benefit to having more exposure to the unexpected object.

Interestingly, there did appear to be a small effect of where the object onset on noticing rates. Subjects were slightly more likely to notice objects that onset at the edge than those that onset near the middle of the display. Given that the task requires monitoring bounces from the edge of the display, this difference in noticing might result from the deliberate allocation of attention to the edges. If so, that pattern provides further support for the idea that noticing happens soon after onset. Even though objects that onset near the middle of the display spend just as much time near the edges as those that onset at the edge, the heightened attention at the edge may be what enhances noticing for the edge-onset objects. No such attentional advantage for the mid-display onset objects may indicate that the window for noticing for objects that onset near the middle of the display might already be closed by the time those objects reach the edge of the display.

This small difference in noticing between edge- and mid-display-onset objects appeared to increase with exposure time. Mid-display-onset noticing rates barely increased with additional exposure time, rising from 41% in the 1.5s condition to 45% in the 5s condition. The edge-onset noticing rates increased by a greater amount, from 41% in the 1.5s condition to 54% in the 5s condition, although this increase is still small relative to the magnitude of the exposure time increase. However, there is a fair amount of variability in the estimate of the difference within each exposure condition; even the largest difference of 8 percentage points between the mid-display and edge onsets in the 5s exposure condition has confidence intervals that include small negative differences. That said, the apparently larger effect in the 2.67s and 5s condition may indicate that the heightened attention to the edge of the display widens the window in which the unexpected object can be detected, and the 1.5s condition is too brief to benefit; alternatively, it may be sufficiently close to the edge to receive the boost in both onset conditions.

This pattern might explain the higher noticing rates observed for slow (longer exposure) than fast (shorter exposure) objects (Kreitz et al., 2016). In that study, subjects monitored a horizontal line through the display and counted how often the attended objects touched the line. The unexpected object traveled horizontally through the display, parallel to this line, staying close to the focus of attention the entire time it was onscreen (including when it first appeared). The increased noticing with longer exposure time in that experiment might result from the unexpected object onsetting near the locus of attention, just as the we observed a greater impact of exposure time when the unexpected object first appeared at the edge of the display. Even in this case, however, the magnitude by which noticing increased was small despite the large increase in exposure time.

If exposure time has minimal impact on subjects' likelihood of noticing an unexpected object, does that imply that subjects in each condition also formed an equally accurate representation of the unexpected object regardless of how long they potentially could inspect it? Since detecting the unexpected object does not guarantee detailed or accurate encoding (Rock & Gutman, 1981; Sagi & Julesz, 1985), might increased exposure time allow noticers to form a more accurate representation of the unexpected object's features? Experiment 3 addresses whether exposure time has an impact on how much detail subjects can encode about the object, even if the overall probability of noticing is relatively unaffected.

EXPERIMENT 3

Experiment 3 relied on a similar procedure to Experiment 1, using two exposure durations—1.5 and 5 seconds—with onset and offset locations near the middle of the display and centered around fixation. Instead of a gray cross, this time the unexpected object was a randomly-chosen color. Rather than a forced-choice identification, subjects were asked to report the object's color with a continuous color slider. This allowed us to collect not just accuracy data, but also precision. Do subjects who have more exposure to the unexpected object also have a better representation of its features? That is, even though additional exposure time does not seem to affect noticing rates substantially, does it affect how much information about the unexpected object noticers are able to extract?

Methods

The materials and preregistration for this experiment are available at https://osf.io/wx5ua/. A demonstration of the task can be found at simonslab.com/mot/temporal_mot_col_demo.html.

Subjects. We aimed to recruit 1000 subjects in order to finish with 250 in each of two conditions, anticipating a 50% exclusion rate. We planned for a larger sample in each condition in order to better evaluate the precision of representations. We recruited 1004 subjects in total.

Materials and procedure. Experiment 3 used the same multiple object tracking task described in the General Methods, but with a mid-gray (#808080) background instead of a light blue one. The fixation square was a 6 x 6 black square with a white, 2 pixel-wide border so that it would remain visible regardless of which color of object passed behind it.

The unexpected object could be onscreen for either 1.5 or 5 seconds. Like Experiment 1 (and unlike Experiment 2), the unexpected object's motion was centered around fixation. It onset from behind an invisible occluder on one side of fixation, passed behind fixation, and exited behind another invisible occluder on the opposite side of fixation, as in Experiment 1 (Figure 4.2).

The unexpected object in this experiment again was a cross, and was one of 12 randomly chosen colors. The hues ranged from 0 to 330 in HSV (hue, saturation, value) space in 30 degree intervals, and the saturation and value were fixed at 50%. We ran two small-sample pilot studies with 41 and 62 subjects prior to the main experiment. The pilot procedure was identical to that of the main experiment but tested different color values for the unexpected object in order to find a set of colors for which noticing was not at ceiling. In the first pilot, we used the same hues with saturation set to 100% and value set to 75%. Noticing was above 90% in both exposure conditions, so we reduced the intensity of the colors and ran the second pilot to verify that the change reduced overall noticing rates before proceeding to the primary experiment. The pilot data are available on the OSF page for this experiment.

Experiment 3 also changed how participants reported the color of the unexpected object. Rather than selecting a color from a drop-down menu of predetermined options, subjects matched the color of the unexpected object with a slider. Clicking and dragging the marker on the slider adjusted the hue of a reference rectangle underneath it. The slider's saturation and value were set to 50%, so it was possible to exactly match the appearance of the unexpected object.

Subjects also completed a digital version of the Farnsworth D-15 task after the question about vision correction (Farnsworth, 1947; code adopted from a digital version of the task by Daniel Flück, https://www.color-blindness.com/color-arrangement-test/). Subjects had to order 15 colored patches by chromaticity by dragging them into empty slots. They were provided with a fixed reference patch and told to complete the task by always selecting the next-most-similar color and dropping it into the next available slot. This task was intended to identify color deficiencies and monitor settings that interfered with the ability to discriminate colors. The procedure of Experiment 3 was otherwise the same as that described in the General Methods.

Results and Discussion

Prior to analysis, we excluded 516 subjects (51% of our sample) according to the criteria described in the General Methods. We retained data from 488 subjects (*n* = 249 in the 1.5s condition and *n* = 239 in the 5s condition). Data are available at https://osf.io/9mz53/.

Noticing. In Experiment 3, we did not classify subjects as noticers and nonnoticers using the criteria from the first two experiments (i.e., reported noticing and correctly identified the shape and/or color). Instead, given that the goal of this study was to examine precision in their representation of the unexpected object, we treated their self-report of noticing as evidence that they saw the unexpected object and then analyzed their responses for the shape, color, and location of the unexpected object conditioned on whether they had reported seeing something new.

Replicating the main pattern from Experiments 1 and 2, self-reported noticing rates were similar across the two exposure conditions (Figure 4.3). 61.5% (95% CI: [55.4, 67.5]) of subjects reported noticing something new in the 1.5s condition, and 64.9% (95% CI: [58.6, 70.7]) of

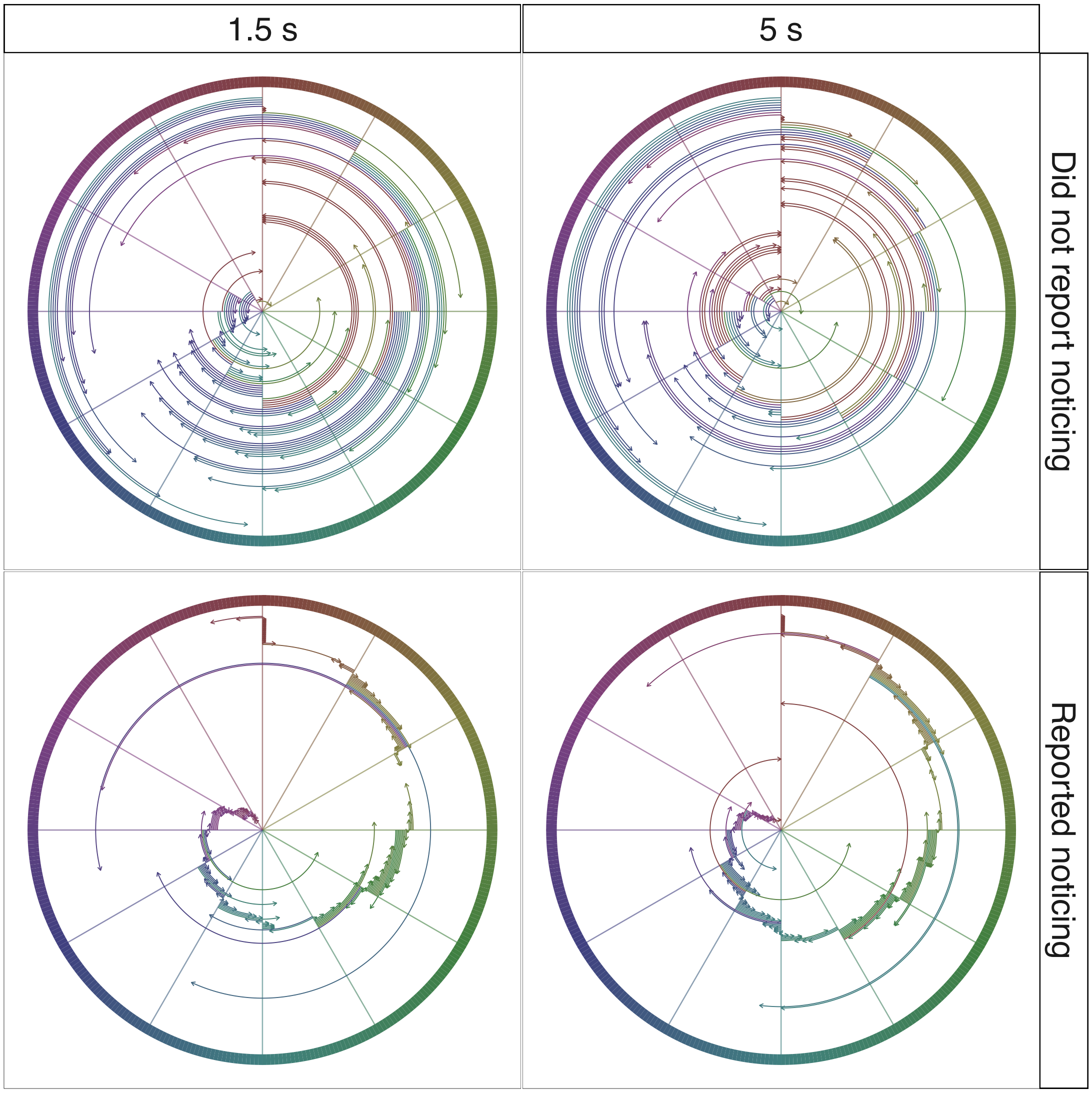

Figure 4.5. Raw color error for each subject. The line segment begins at the true hue of the unexpected object, indicated by the spokes in the wheel. The line ends with the arrow pointing to the subjects' reported hue, and the line segment is colored to match the reported hue.

subjects reported noticing something new in the 5s condition for a difference of -3.4 percentage points (95% CI: [-12.0, 5.1]).

Feature reporting accuracy. Self-reported noticers in both exposure conditions were highly accurate at reporting the shape of the unexpected object. 87.6% (95% CI: [82.4, 92.8]) of the self-reported noticers in the 1.5s exposure condition correctly reported the shape of the unexpected object, and 88.4% (95% CI: [83.2, 92.9]) of them did so in the 5s condition. Non-noticers in the 1.5s and 5s condition were approximately at chance levels (11.1%) in selecting the unexpected object's shape (1.5s: 10.4%, 95% CI: [4.2, 16.7]; 5s: 8.3%, 95% CI: [2.4, 14.3]).

Self-reported noticers in both conditions were extremely accurate at reporting the color, with a circular mean error of -1.1 degrees (angular deviation = 27.2) in the 1.5s exposure condition and a circular mean error of -1.9 degrees (angular deviation = 29.2) in the 5s condition (Figure 4.5). The difference in means between the two conditions was .83 degrees (95% CI: [-4.3, 5.4]), and the ratio of the circular variance between the two conditions was .87 (95% CI: [.45, 1.66]). Thus the two conditions appear to be not just equally accurate at reporting the color, but equally precise. Non-noticers, in contrast, had a mean circular error of -109.9 degrees (angular deviation = 79.9) in the 1.5s condition and 171.8 (angular deviation = 79.3) in the 5s condition (Figure 4.5). If we generate 1000 samples of 100 subjects selecting a color completely randomly from any point on the color wheel, the mean angular deviation of the error would be 77.4 degrees (95% CI: [72.6, 80.4]). In other words, the variability in color selection shown by non-noticers is what we would expect from chance responding.

Location reports. Consistent with the results of Experiments 1 and 2, self-reported noticers localized the unexpected object near its onset location on average (Figure 4.4). Noticers had a tighter spread to their vertical placements than did non-noticers, as in Experiments 1 and 2, placing their objects at at average of 286 pixels *(SD* = 75) compared to 295 (*SD* = 159) for non-noticers. Noticers placed the object on the onset side of fixation 70% of the time; non-noticers did so 45% of the time. Noticers generally placed the object closest to the onset location, whereas non-noticer localizations were more evenly distributed (Table 4.4). The tendency to localize the unexpected object near the onset location is easier to see in the 5s condition where there is more space between onset and fixation than in the 1.5s condition, when fixation, onset, and offset were more compressed.

Experiment 3 closely replicates the pattern of results observed in the previous experiments. First, dramatically reduced exposure time to the unexpected object had comparatively little effect on the probability that subjects would notice it. Second, subjects tended to report noticing the unexpected object near its onset location. An appreciable proportion of noticing events apparently occur soon after the unexpected object appears.

Noticed Unexpected Object	Time Onscreen	Unexpected Velocity	Distance to Fixation	Distance to Onset	Distance to Offset
No	1.5	Right to left	234.7 (148.7)	246.4 (151.4)	247.0 (139.9)
No	1.5	Left to right	199.0 (148.8)	214.0 (137.1)	215.9 (148.3)
No	5.0	Right to left	174.6 (132.2)	307.2 (160.3)	338.9 (134.0)
No	5.0	Left to right	223.7 (139.8)	352.0 (135.7)	339.4 (197.2)
Yes	1.5	Right to left	84.1 (89.6)	106.1 (90.7)	115.3 (85.0)
Yes	1.5	Left to right	81.7 (103.8)	104.6 (93.8)	117.8 (105.3)
Yes	5.0	Right to left	218.9 (108.4)	157.1 (136.7)	459.3 (148.2)
Yes	5.0	Left to right	244.1 (112.6)	144.7 (126.2)	492.7 (148.0)

Table 4.4. Average Euclidean distance, in pixels, of subjects' location reports for the unexpected object to onset, fixation, and offset in Experiment 3. For each condition and direction of motion, we averaged the distance between each individual location report and the location of onset in that condition, fixation, and the location for offset in that condition.

The results of Experiment 3 further suggest that not only is there minimal impact of exposure time on the likelihood of noticing the unexpected object, but also minimal impact on the precision of the representation of that object. Subjects who noticed the unexpected object were equally accurate at reporting the object's shape and color whether they were exposed to it for 1.5 seconds or 5 seconds. The extra exposure time did not improve accuracy or precision. Indeed, there was almost no room to improve: The circular mean of subjects' reports of the color in the 1.5s condition were already nearly perfect. However, the accuracy for reporting the shape, which was not at ceiling and could have reflected a difference, was also the same across the two exposure conditions. It appears that subjects in both conditions had equally accurate representations of the unexpected object despite a large difference in exposure time. In contrast, non-noticers were inaccurate in both their color and shape reports. Although we did not observe a benefit of exposure time on feature reporting accuracy, our unexpected object had only two features (color and shape), and it is possible that such a benefit would emerge with a more complex object.

GENERAL DISCUSSION

Across three experiments, substantially reducing the amount of time an unexpected object was onscreen, by 50% or even 70%, had only a modest impact on the proportion of subjects noticing it. The largest difference in noticing as a function of exposure time was for unexpected objects that onset from the edge of the display in Experiment 2, in which there was a 12.7 percentage point difference between the 5 second and 1.5 second exposure durations.

The window for noticing an unexpected object appears to be brief relative to the amount of time it is visible onscreen. Even though subjects have more opportunity to detect the unexpected object the longer it remains onscreen, the vast majority of noticing events occur in the first 1.5 seconds or not at all. This pattern, replicated in all three experiments, indicates that unexpected objects are not noticed as a result of a gradual accumulation of signal across the entire exposure duration, but more as a result of a rapid process concentrated early in the unexpected object’s lifespan.

The results from subjects' reports of when they first noticed the unexpected object help narrow down when these noticing events occur. Subjects who noticed the unexpected object tended to report first seeing it near the onset location in all conditions. The onset event itself might trigger noticing, with a brief window of opportunity for detection shutting rapidly afterward. The abrupt appearance of a new object in the absence of other events or distractions can provide a strong attention signal in other tasks (e.g., Yantis & Jonides, 1984). However, the effectiveness of such onsets in capturing attention is reduced when they coincide with other dynamic events (Rensink, 2000).

Alternatively, noticing might not be triggered by the onset itself, but there may nevertheless be a brief period during which a new object in the display can be detected before it is filtered out by attentional selection. The noticing data and location reports do show that offset events do not seem to trigger noticing.

Experiment 3 further revealed that not only does exposure time have little impact on noticing, it also seems to have little impact on the quality of the representation of the unexpected object's features. This presents the counterintuitive possibility that if an object is noticed, it is also represented at the highest precision it could be. More exposure to the object after the initial noticing event did not further improve the representation. If so, detection and representation of the unexpected object might be an all-or-nothing process. It either happens in its entirety within a short window after onset, or not at all. Alternatively, detection and representation may be distinct from each other, but on a timescale shorter than 1.5 seconds.

What precisely unfolds in those initial 1.5 seconds remains an avenue for exploration. We can be confident that the early noticing we observed is not due to a ceiling effect. Noticing rates hovered around 50%, meaning that the unexpected object was not so conspicuous that everyone noticed it. Many people missed the object, but those who noticed it appeared to do so rapidly. The onset event itself, and the associated visual change, might trigger noticing of the unexpected object. Or, once an additional object appears, the attentional system might rapidly accumulate signal which passes some threshold for detection with a given probability (meaning it is noticed) or does not (meaning it is missed). There may even be competition between the signal generated by the new object and suppression of task-irrelevant information, and this suppression of irrelevant information could affect the likelihood of detecting the object after onset.

The time intervals used in the three experiments and the reports of the unexpected object's location are not precise enough to definitively support one of these models. However, these results can rule out other possible accounts of how noticing unfolds in sustained inattentional blindness. Noticing was not triggered by offset events or by the unexpected object crossing fixation, and noticing is neither a slow process that unfolds over time, nor a slow accumulation of evidence that accelerates the longer the unexpected object is visible. Rather, noticing in these sustained inattentional blindness tasks may largely be a process that happens almost immediately or not at all.

CONSTRANTS ON GENERALITY

We expect these results to generalize to similar tasks and to be robust to arbitrary choices about the appearance and behavior of the stimuli and display, provided that overall noticing levels are not driven to ceiling or floor. Other studies have used sustained inattentional blindness tasks like these across a range of settings (laboratory, public, online), so we expect our results to generalize to adult populations with normal or correct-to-normal visual acuity in both online and in-person settings.

Our results in Experiment 2 suggest that differences in noticing rates as a function of spatial attention likely would not interact with the exposure duration effects we examined. Similarity effects, wherein objects similar to the attended set tend to be noticed at high rates whereas objects in the ignored set tend to be noticed at much lower rates (e.g. Most et al., 2001) might interact with the effects of exposure duration we observed in these experiments.

Precision for color selection was near ceiling in Experiment 3, but that does not necessarily mean that exposure time has no effect on how accurately subjects can report features. Perhaps the task used in Experiment 3 was too easy for subjects, and a more difficult task would reveal effects of exposure; perhaps a more complex object with more features than shape and color would benefit from additional time, even though the simple object we used did not. Additionally, it is possible that much longer exposures would yield more advantage and that 5s not enough additional time to reveal a benefit of additional exposure. Exposure time has an impact on encoding and representation in other tasks, such as change blindness (Brady, Konkle, Oliva, & Alvarez, 2015). These effects may still be present in inattentional blindness as well, but may either saturate at 1.5 seconds, require a harder task, require even more time to reveal themselves, or some combination of these factors.

Our results might also be limited to simple tasks and displays and might not generalize to more real-world or video-based tasks where the content and interaction between the action of the video and the unexpected object can affect the time course of noticing (Becklen & Cervone, 1983).

CHAPTER 5: CONCLUSION

The sense that we experience the world in a seamless, detailed stream of information is an illusion. The strict capacity limits at every level of our visual system preclude this possibility, from the neuron upward. We cope by being selective about the information we process, and filtering unwanted information. There is much more of the latter than the former, and as inattentional blindness reveals, this means that our experience of the world is shaped as much by what we don't see as by what we do.

The results of the experiments described in Chapter 2 reveal the nonselective nature of suppression. When the selective attention system is presented with a set of distractors that are heterogeneous with respect to the critical feature, it responds by inhibiting any feature value other than the attended one. When subjects ignored a set containing distractors of several colors or ignored a color that changed on each trial, they unsurprisingly failed to notice unexpected objects that shared a color with any ignored object. More strikingly, however, they also failed to notice an unexpected object bearing a completely novel color. It was noticed no more frequently than unexpected objects that shared a color that was actively being ignored. Attending to the variable color did not yield the same results; only unexpected objects that shared the attended color were noticed at any appreciable rate. This suggests an asymmetry in selective attention: We select narrowly and inhibit broadly. This aligns well with the similar asymmetry in the amount of information we do not want to select versus the amount we do any time we engage selective attention.

Further, these results undermine the notion that we select desired information without actively inhibiting unwanted information. When subjects ignored the variable colors, their noticing rates for the novel object were exceedingly low, and were similar to the noticing rates for the color being ignored on the current trial. If this were due simply to subjects not picking up the information, rather than actively inhibiting it, then the noticing rates for the novel-colored object ought to be the same when subjects were attending to the variable color. After all, the novel color is just as task-irrelevant in either case, and subjects have no reason to select it. However, we observe substantially higher noticing rates for the novel-colored object when subjects are attending to the variable color. This is not a mere novelty effect, because this identical object goes utterly unnoticed when subjects are ignoring color variability. Rather, it seems that when subjects must attend to variable features, the novel color does not get suppressed as aggressively and so reaches awareness at higher rates. No “selection-only” model can account for this difference; only suppression can explain why it is almost completely unnoticed in one case but noticed much more frequently in the other.

These results pose a stiff challenge to models that propose that attentional selection is the more critical factor in noticing, or that no filtering mechanism is needed at all to explain why we fail to notice unexpected objects (Neisser, 1976; Neisser, 1979). Such models were difficult to refute with homogenous sets of objects, since either a selection or inhibition account could explain the results. Results of these attention-set experiments were sometimes contextualized in this framework (e.g. Most et al., 2005), and studies that tried to untangle the mechanisms did not have displays and unexpected objects that yielded different predictions for the two alternatives (e.g. Goldstein & Beck, 2016). However, with the displays and unexpected objects employed in Chapter 2, the effect of suppression was demonstrated independently of that of selection, and suppression was shown to act more broadly than selection does.

Chapter 3 revealed similarly broad effects of suppression over space. Subjects performed a task in a large, busy display, only a portion of which was truly relevant to the task of navigating back and forth while avoiding obstacles. The allocation of attention followed intuitively from this goal, narrowing around the avatar subjects controlled and the area that corresponded to the highest risk of colliding with a hazard.

While these experiments are far from the first to explore how goals influence attentional allocation, or how attention varies according to environment, the inattentional blindness paradigm offers a particular advantage: namely, subjects' attention can be measured without the need for a secondary task. They do not need to divide their attention between the task meant to influence their attention and the task meant to measure it. This is in contrast to some prior work on the effect of task on attention, such as one that required subjects to either drive a virtual car down a road or simply watch a video as if they were a passenger in said car (Wallis & Büllthoff, 2000). Subjects' attention was measured by their performance on a secondary change-detection task featuring objects sitting on the side of the road. It is entirely possible, in divided-attention contexts such as these, that subject behavior on the primary task changes as a result of having to dedicate attention to a secondary task; this may in turn have affected how attention was allocated, particularly since one condition was essentially just a change detection task and the other was a dual task. This unequal cognitive load further limits the conclusions that can be drawn about the effects on attention. Using inattentional blindness as the measure of attention circumvents the problems that arise from these dual-task designs.

The series of experiments described in Chapter 4 provide an indication of how unexpected objects break through suppression in sustained inattentional blindness. Across three experiments, large increases in exposure time increased subjects' noticing rates only slightly. Subjects' location reports indicate that this noticing rate is so insensitive to exposure time because the unexpected object tends to be noticed soon after it appears. The probability of noticing is not constant the entire time the unexpected object is onscreen, nor does it appear to break through at random. It instead appears that the onset of the unexpected object triggers a narrow window of time in which it can be noticed.

The time course suggested by these results speaks to how thoroughly task-irrelevant information is suppressed during these tasks. It's not simply that suppression makes subjects unlikely to notice the object overall; rather, it drives the probability of noticing the object down virtually to zero after a certain point. Noticing is not necessarily inevitable given enough time. This result rules out a wide array of accounts of noticing that require an extended amount of time, or at iterative process in which the items in the display are continuously re-evaluated as attention shifts (Neisser, 1976). These results, particularly in light of the findings that subjects are equally accurate at reporting object features, also complicate the proposition that noticing might unfold in stages, or that we might build up a representation over time. Instead, it seems noticing is closer to an all-or-nothing process that occurs soon after onset.

Inattentional blindness is frequently described as a "failure of awareness." The results described across these three chapters indicate that there is more to inattentional blindness than a failure of attentional selection, or a failure to orient to a new object. Looked at another way, inattentional blindness demonstrates our incredible ability to filter unwanted information. When we want to concentrate and narrow our focus to one small piece of a busy world, our suppressive capabilities can erase distracting objects from awareness entirely—even odd or novel events that we might want to become aware of.

If an object does not break through immediately, it is unlikely to break through in the following few seconds. When the distracting information is heterogeneous with respect to its features, any feature other than the attended one is suppressed, regardless of whether or not it is present in a display. When the distracting information is scattered widely across space, only the task-relevant areas are spared. Suppression acts quickly and broadly to block irrelevant information from awareness. This is likely the state we find ourselves in more often than we realize; concentrating, and oblivious, never to realize the scope of the information we did not encode. The world is not ours at a glance, teeming with detail and richness. Instead we sit in the dark, the better to illuminate the small piece of the world we can take in moment to moment.

REFERENCES

Adamo, M., Pun, C., Pratt, J., & Ferber, S. (2008). Your divided attention, please! The maintenance of multiple attentional control sets over distinct regions in space. *Cognition, 107*(1), 295–303. https://doi.org/10.1016/j.cognition.2007.07.003

Agostinelli, C., & Lund, U. (2017). R package 'circular': Circular Statistics (version 0.4-93). URL https://r-forge.r-project.org/projects/circular/

Aimola Davies, A. M., Waterman, S., White, R. C., & Davies, M. (2013). When you fail to see what you were told to look for: Inattentional blindness and task instructions. *Consciousness and Cognition, 22*(1), 221–230. http://doi.org/10.1016/j.concog.2012.11.015

Alvarez, G. a, & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15*(2), 106–111. https://doi.org/10.1167/2.7.273

Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological Science, 16*(8), 637–643. https://doi.org/10.1111/j.1467-9280.2005.01587.x

Bacon, W. F., & Egeth, H. E. (1994). Overriding stimulus-driven attentional capture. *Perception & Psychophysics, 55*(5), 485-496.

Beanland, V., & Pammer, K. (2010a). Looking without seeing or seeing without looking? Eye movements in sustained inattentional blindness. *Vision Research, 50*, 977–988. http://doi.org/10.1016/j.visres.2010.02.024

Beanland, V., & Pammer, K. (2010b). Gorilla watching: Effects of exposure and expectations on inattentional blindness. In W. Christensen, E. Schier, and J. Sutton (Eds.), *ASCS09:*

Proceedings of the 9th Conference of the Australian Society for Cognitive Science (pp. 12-20). Sydney: Macquarie Centre for Cognitive Science.

Becklen, R. & Cervone, D. (1983). Selective looking and the noticing of unexpected events. *Memory & Cognition, 11*(6), 601-608.

Bichot, N. P., Rossi, A. F., & Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area 4. *Science, 308,* 529–534.

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

Brady, T. F., Konkle, T., Oliva, A., & Alvarez, G. A. (2015). Detecting changes in real-world objects: The relationship between visual long-term memory and change blindness. *Communicative & Integrative Biology, 2*(1), 8–11. https://doi.org/10.4161/cib.2.1.7297

Brawn, P., & Snowden, R. J. (1999). Can one pay attention to a particular color? *Perception & Psychophysics, 61*(5), 860-873.

Bredemeier, K., & Simons, D. J. (2012). Working memory and inattentional blindness. *Psychonomic Bulletin & Review, 19*(2), 239–244. http://doi.org/10.3758/s13423-011-0204-8

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research, 51*(13), 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience, 7*(3), 308–313. https://doi.org/10.1038/nn1194

Cartwright-Finch, U., & Lavie, N. (2007). The role of perceptual load in inattentional blindness. *Cognition, 102*(3), 321–340. http://doi.org/10.1016/j.cognition.2006.01.002

Chabris, C. F., Weinberger, A., Fontaine, M., & Simons, D. J. (2011). You do not talk about Fight Club if you do not notice Fight Club: Inattentional blindness for a simulated real-world assault. *I-Perception, 2*, 150–153.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and two ears. *Journal of the Acoustical Society of America, 25*, 975-979.

Clement, A., Stothart, C., Drew, T., & Brockmole, J. R. (in press). Semantic associations do not modulate the visual awareness of objects. *Quarterly Journal of Experimental Psychology.*

Cornell, A. D. (1959). Further experiments in apparitional observation. *Journal of the Society for Psychical Research, 40*(706), 409–418.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective attention. *Annual Review of Neuroscience, 18*, 193–222.

Drew, T., & Stothart, C. (2016). Clarifying the role of target similarity, task relevance and feature-based suppression during sustained inattentional blindness. J*ournal of Vision, 16*(15), 1–9. http://doi.org/10.1167/16.15.13.doi

Drew, T., Võ, M. L., & Wolfe, J. M. (2013). The Invisible Gorilla Strikes Again : Sustained Inattentional Blindness in Expert Observers. *Psychological Science, 24*(9), 18–185348. http://doi.org/10.1177/0956797613479386

Efron, B., & Tibshirani, R. J. (1993). An introduction to the bootstrap. CRC press.

Egeth, H. E., & Yantis, S. (1997). Visual attention: control, representation, and time course. *Annual Review of Psychology, 48*, 269–297. https://doi.org/10.1146/annurev.psych.48.1.269

Farnsworth, D. (1947). The Farnsworth dichotomous test for color blindness: Panel D-15 manual. New York: The Psychological Corporation.

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 18*(4), 1030.

Garnier, S. (2018). viridis: Default Color Maps from 'matplotlib'. R package version 0.5.1. https://CRAN.R-project.org/package=viridis

Goldin, G., Darlow, A. (2013). TurkGate (Version 0.4.0) [Software]. Available from http://gideongoldin.github.com/TurkGate/

Goldstein, R. R., & Beck, M. R. (2016). Inattentional blindness: A combination of a relational set and a feature inhibition set? *Attention, Perception, & Psychophysics, 78*, 1245–1254. http://doi.org/10.3758/s13414-016-1091-x

Grimes, J. (1996) On the failure to detect changes in scenes across saccades. In *Perception (Vancouver Studies in Cognitive Science)* (Vol. 5) (Akins, K., ed.), pp. 89–110, Oxford University Press.

Haines, R.F. (1991). A breakdown in simultaneous information processing. In G. Obrecht & L.W. Stark (Eds.), *Presbyopia Research* (pp. 171–175). New York: Plenum Press.

Henry, L. & Wickham, H. (2017). purrr: Functional programming tools. R package version 0.2.4. https://CRAN.R-project.org/package=purrr

Hollingworth, A. (2003). Failures of retrieval and comparison constrain change detection in natural scenes. J*ournal of Experimental Psychology: Human Perception and Performance, 29*(2), 388–403. https://doi.org/10.1037/0096-1523.29.2.388

Hyman, Ira E, J., Boss, S. M., Wise, B. M., McKenzie, K. E., & Caggiano, J. M. (2010). Did you see the unicycling clown? Inattentional blindness while walking and talking on a cell phone. *Applied Cognitive Psychology, 24*, 597–607. http://doi.org/10.1002/acp

Ishihara, S. (1990). Ishihara's test for color blindness: 38 Plates Edition. Tokyo, Japan: Kanehara and Co., Ltd.

Koivisto, M., Hyönä, J., & Revonsuo, A. (2004). The effects of eye movements, spatial attention, and stimulus features on inattentional blindness. *Vision Research, 44*(27), 3211–3221. http://doi.org/10.1016/j.visres.2004.07.026

Kreitz, C., Furley, P., & Memmert, D. (2016). Inattentional blindness is influenced by exposure time not motion speed. *Quarterly Journal of Experimental Psychology, 69*(3), 495-505. http://dx.doi.org/10.1080/17470218.2015.1055771

Kreitz, C., Furley, P., Memmert, D., & Simons, D. J. (2015). Inattentional blindness and individual differences in cognitive abilities. *PLoS ONE, 10*(8), 1–27. http://doi.org/10.1371/journal.pone.0134675

Kuhn, G., & Rensink, R. A. (2016). The vanishing ball illusion: A new perspective on the perception of dynamic events. *Cognition, 148*, 64–70. https://doi.org/10.1016/j.cognition.2015.12.003

Liu, T., Stevens, S. T., & Carrasco, M. (2007). Comparing the time course and efficacy of spatial and feature-based attention. *Vision Research, 47*(1), 108–113. https://doi.org/10.1016/j.visres.2006.09.017

Mack, A., & Rock, I. (1998). Inattentional blindness. MIT press.

Mack, A., Tang, B., Tuma, R., Kahn, S., & Rock, I. (1992). Perceptual organization and attention. *Cognitive Psychology, 24*, 475–501.

Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences, 9*(6), 296-305. https://doi.org/10.1016/j.tics.2005.04.010

McAdams, C. J., & Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience, 19*(1), 431-441.

Memmert, D. (2006). The effects of eye movements, age, and expertise on inattentional blindness. *Consciousness and Cognition, 15*(3), 620–627. http://doi.org/10.1016/j.concog.2006.01.001

Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics, 66*(8), 1268–1281. https://doi.org/10.3758/BF03194997

Moher, J., Lakshmanan, B. M., Egeth, H. E., & Ewen, J. B. (2014). Inhibition drives early feature-based attention. *Psychological Science, 25*(2), 315–324. https://doi.org/10.1177/0956797613511257

Moore, C. M., & Egeth, H. (1997). Perception without attention: evidence of grouping under conditions of inattention. *Journal of Experimental Psychology. Human Perception and Performance, 23*(2), 339–352. https://doi.org/10.1037/0096-1523.23.2.339

Moore, C. M., & Egeth, H. (1998). How does feature-based attention affect visual processing? *Journal of Experimental Psychology: Human Perception and Performance, 24*(4), 1296-1310.

Most, S. B. (2013). Setting sights higher: Category-level attentional set modulates sustained inattentional blindness. *Psychological Research, 77*(2), 139–146. http://doi.org/10.1007/s00426-011-0379-7

Most, S. B., Simons, D. J., Scholl, B. J., & Chabris, C. F. (2000). Sustained inattentional blindness: The role of location in the detection of unexpected dynamic events. *Psyche, 6*(14), 1059–1074.

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*(1), 9–17. http://doi.org/10.1111/1467-9280.00303

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*(1), 217–242. http://doi.org/10.1037/0033-295X.112.1.217

Müller, N. G., Mollenhaur, M., Rösler, A., & Kleinschmidt, A. (2005). The attentional field has a Mexican hat distribution. *Vision Research, 45,* 1129-1137.

Nakajima, M., Schmitt, L. I., & Halassa, M. M. (2019). Prefrontal cortex regulates sensory filtering through a basal ganglia-to-thalamus pathway. *Neuron, 103,* 1-14. https://doi.org/10.1016/j.neuron.2019.05.026

Neisser, U. (1976). Cognition and reality: Principles and implications of cognitive psychology. San Francisco: W. H. Freeman.

Neisser, U. (1979). The control of information pickup in selective looking. In A. D. Pick (Ed.), *Perception and its development: A tribute to Eleanor J. Gibson* (pp. 201–219). Hillsdale, NJ: Erlbaum.

Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually specified events. Cognitive Psychology, 7(4), 480–494. http://doi.org/10.1016/0010-0285(75)90019-5

Neisser, U., & Dube, E. F. (1978). Interrupting the perceptual cycle: When do we notice unexpected events? Paper presented at the meeting of the Eastern Psychological Association, Washington, D.C., March 1978.

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

Newby, E. A., & Rock, I. (2001). Location and attention. *Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology, 54*(1), 155–168. https://doi.org/10.1080/02724980042000066

O'Connor, D. H., Fukui, M. M., Pinsk, M. A., & Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience, 5*(11), 1203-1209. doi:10.1038/nn957

Pedersen, T. L. (2019). ggforce: Accelerating 'ggplot2'. R package version 0.2.2. https://CRAN.R-project.org/package=ggforce

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 3*2(1), 3-25.

Posner, M. I., Snyder, C. R. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General, 109*(2), 160-174.

R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18*(3), 849.

Rees, G., Russell, C., Frith, C. D., & Driver, J. (1999). Inattentional blindness versus inattentional amnesia for fixated but ignored words. *Science, 286,* 2504-2507. 10.1126/science.286.5449.2504

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 37*(3), 1–6. https://doi.org/10.1111/j.1467-9280.1997.tb00427.x

Rock, I. & Gutman, D. (1981). The effect of inattention on form perception. *Journal of Experimental Psychology: Human Perception and Performance, 7*(2), 275 - 285.

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature neuroscience, 5*(7), 631.

Sagi, D. & Julesz, B. (1985). Detection versus discrimination of visual orientation. *Perception, 14*, 619 - 628.

Saiki, J., & Holcombe, A. O. (2012). Blindness to a simultaneous change of all elements in a scene, unless there is a change in summary statistics. *Journal of Vision, 12*(3):2, 1–11. https://doi.org/10.1167/12.3.2

Scholl, B. J. (2009). What have we learned about attention from multiple object tracking (and vice versa). *Computation, cognition, and Pylyshyn*, 49-78.

Shih, S., & Sperling, G. (1996). Is there feature-based selection in visual search? *Journal of Experimental Psychology: Human Perception and Performance, 22*(3), 758-779.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science, 22*(11), 1359-1366. http://doi.org/10.1177/0956797611417632

Simons, D. J. (2010). Monkeying around with the gorillas in our midst: Familiarity with an inattentional-blindness task does not improve the detection of unexpected events. *I-Perception, 1*(1), 3–6. http://doi.org/10.1068/i0386

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28*(9), 1059–1074. http://doi.org/10.1068/p2952

Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review, 5*(4), 644–649. https://doi.org/10.3758/BF03208840

Simons, D. J., & Jensen, M. S. (2009). The effects of individual differences and task difficulty on inattentional blindness. *Psychonomic Bulletin & Review, 16*(2), 398–403. http://doi.org/10.3758/PBR.16.2.398

Simons, D. J., & Schlosser, M. D. (2017). Inattentional blindness for a gun during a simulated police vehicle stop. *Cognitive Research: Principles and Implications, 2*(37). http://doi.org/10.1186/s41235-017-0074-3

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied, 74*(11), 1–29. doi:10.1037/h0093759

Spitzer, H., Desimone, R., & Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science, 240*(4850), 338-340.

Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology, 24*(17), 1985–1988. https://doi.org/10.1016/j.cub.2014.07.030

Stothart, C. R., Boot, W. R., & Simons, D. J. (2015). Using Mechanical Turk to assess the effects of age and spatial proximity on inattentional blindness. *Collabra, 1*(1), 1–7.

Stothart, C. R., Wright, T. J., Simons, D. J., & Boot, W. R. (2017). The costs (or benefits) associated with attended objects do little to influence inattentional blindness. *Acta Psychologica*, *173*, 101-105.

Stothart, C., Simons, D. J., Boot, W. R., & Wright, T. J. (2019). What to where: The right attention set for the wrong location. *Perception, 0*(0), 1–14. https://doi.org/10.1177/0301006619854302

Theeuwes, J. (1989). Effects of location and form cuing on the allocation of attention in the visual field. *Acta Psychologica, 72,* 177-192.

Theeuwes, J. (2004). Top-down search strategies cannot override attentional capture. *Psychonomic Bulletin and Review, 11*(1), 65–70. https://doi.org/10.3758/BF03206462

Treisman, A. M. (1964a). The effect of irrelevant material on the efficiency of selective listening. *American Journal of Psychology, 77*, 533-546.

Treisman, A. M. (1964b). Verbal cues, language, and meaning in selective attention. *American Journal of Psychology, 77*, 206- 219.

Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General, 114*(3), 285-310. doi:http://dx.doi.org/10.1037/0096-3445.114.3.285

Treue, S., & Trujillo, J. C. M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399,* 575-579.

Tse, P. U. (2005). Voluntary attention modulates the brightness of overlapping transparent surfaces. *Vision Research, 45*(9), 1095–1098

Underwood, G., Chapman, P., Brocklehurst, N., Underwood, J., & Crundall, D. (2003). Visual attention while driving: Sequences of eye fixations made by experienced and novice drivers. *Ergonomics, 46,* 629-646.

Vingrys, A. J., & King-smith, P. E. (1988). A quantitative scoring technique for panel tests of color vision. *Investigative Ophthamology & Visual Science, 29*(1), 50–63.

Wallis, G., & Bulthoff, H. (2000). What's scene and not seen: Influences of movement and task upon what we see. *Visual Cognition, 7*, 175–190.

Ward, E. J., & Scholl, B. J. (2015). Inattentional blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness. *Psychonomic Bulletin & Review, 22*(3), 722–727. http://doi.org/10.3758/s13423-014-0745-8

Wickham, H. *ggplot2: Elegant Graphics for Data Analysis.* Springer-Verlag New York, 2009.

Wickham, H. (2018). stringr: Simple, consistent wrappers for common string operations. R package version 1.3.0. https://CRAN.R-project.org/package=stringr

Wickham, H., François, R., Henry, H., & Müller, K. (2017). dplyr: A grammar of data manipulation. R package version 0.7.4. https://CRAN.R-project.org/package=dplyr

Wickham, H. & Henry, L. (2018). tidyr: Easily tidy data with 'spread()' and 'gather()' functions. R package version 0.8.0. https://CRAN.R-project.org/package=tidyr

Wolfe, J. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review, 1*(2), 202 - 238.

Wood, N. L., & Cowan, N. (1995). The Cocktail Party Phenomenon Revisited: Attention and Memory in the Classic Selective Listening Procedure of Cherry (1953). *Journal of Experimental Psychology: General, 124*(3), 243–262. https://doi.org/10.1037/0096-3445.124.3.243

Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: evidence from visual search. Journal of Experimental Psychology: Human Perception and Performance, 10(5), 601.