Inattentional Blindness in a Coupled Perceptual–Cognitive System

Will Bridewell (will.bridewell@nrl.navy.mil)
Paul F. Bello (paul.bello@nrl.navy.mil)
U.S. Naval Research Laboratory
Washington, DC 20375 USA

Abstract

Attention is thought to be a part of a larger cluster of mechanisms that serve to orient a cognitive system, to filter contents with respect to their task relevance, and to devote more computation to certain options than to others. All these activities proceed under the plausible assumption that not all information can be or ought to be processed for a system to satisfice in an ever changing world. In this paper, we describe an attention-centric cognitive system called ARCADIA that demonstrates the orienting, filtering, and resource-skewing functions mentioned above. The demonstration involves maintaining focus on cognitive tasks in a dynamic environment. While ARCADIA carries out a task, limits on its attentional capacity result in "inattentional blindness" under circumstances analogous to those where people fail to perceive otherwise salient stimuli.

Keywords: attention; perception; vision; cognitive model; inattentional blindness

Introduction

Imagine walking across a familiar courtyard on a clear day, your cellphone rings, and you answer. Your friend invites you to a café, and as you set a time to meet, a stranger with a clipboard asks whether you noticed anything out of place. You say, "No," and look around briefly. The stranger asks, more pointedly, if you saw the unicycling clown. You did not, despite the facts that the sun is shining, that you were aware enough to avoid stumbling into anyone, that the clown was wearing bright clothes, and that he unicycled in your general direction. According to Hyman and colleagues (2010), this scenario is common with only 25% of people talking on cell phones noticing the clown compared to over 50% of others strolling through the courtyard witnessing him clearly. Their research was another poignant example of *inattentional blindness* (Mack & Rock, 1998), joining the ranks of "invisible gorilla" studies (Drew, Vo, & Wolfe, 2013; Simons & Chabris, 1999) that demonstrate the human inability to notice otherwise salient stimuli when cognitively preoccupied.

Research on inattentional blindness suggests that people consciously experience only those objects and events that receive attention. Consequently, failing to attend to an object entails failing to explicitly perceive a stimulus and subsequently failing to verbally report on it. Keeping in mind that attention is widely considered to be extremely limited, strategic focus on one aspect of the environment (e.g., a phone conversation) leaves the perceiver open to missing much of what is going on around them (e.g., a unicycling clown). Perhaps surprisingly, inattentional blindness occurs even when preattentively processed, salient, visual features (e.g., contrast in color, orientation, luminance, motion) would pop out during visual search (Yantis & Egeth, 1999). This finding indicates that saliency-based models of visual attention (Borji & Itti, 2013; Itti, Koch, & Niebur, 1998) are incomplete.

In addition to this incompleteness, what does inattentional blindness say about the character of attention and the models that researchers build? Specifically, the phenomenon indicates that for an object representation to enter awareness, to enter working memory, to be reportable, the representation must receive attention. Unsurprisingly, there are disagreements on this point. Wolfe (1999) attributes inattentional blindness to a failure of memory and says that we process and see everything in our visual field but then forget some of it, possibly due to lack of attention. That is, we are aware of everything, but we may not remember or report on the information. In contrast, Mack (2008) says that although we process and encode everything in our visual field, we only explicitly perceive those objects that receive attention. Her story is supported by recent experiments that appear to rule out memory failures (Ward & Scholl, 2015). In either case, this distinction depends on a notion of seeing that invokes conscious awareness, a topic that continues to elude computational modeling.

With consciousness out of the picture, what characteristics of attention should a cognitive model address? Studies on inattentional blindness tell us that

- cognitive tasks can interfere with the perception of visually presented objects,
- explicit object representations that do not receive attention do not persist,
- the properties (shape, color, etc.) of representations that are not persistent cannot be reported,
- visual objects that do not receive attention are still processed and represented.

Most importantly, an accurate model of attention must account for cases where otherwise salient objects go unseen.

Answering these questions requires a modeling approach that can integrate cognition and perception, where attention plays a role in memory, and where there is a distinction between preattentive and attentive information processing. One cognitive system that meets these requirements is ARCADIA.[1] After briefly describing the system, we discuss a visual processing model and highlight how it combines top-down (i.e., cognitive) and bottom-up (i.e., perceptual) factors to construct and remember object representations. We then provide an overview of early inattentional blindness experiments reported by Mack and Rock (1998). Finally, we present

[1] **A**daptive, **R**eflective **C**ognition in an **A**ttention-**D**riven **I**ntegrated **A**rchitecture

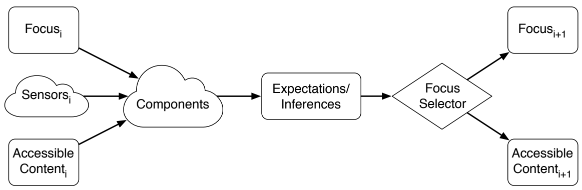

Figure 1: ARCADIA's cognitive cycle.

a model in ARCADIA that carries out a cognitive task and show that it accounts for instances of inattentional blindness by virtue of sharing the four characteristics mentioned above.

ARCADIA's Architectural Features

ARCADIA embodies a theoretical commitment to attention as a central, integrating mechanism for perception, cognition, and action. The system operates in distinct cognitive cycles that are guided by a *focus of attention*. ARCADIA's informational structure is tripartite, consisting of *components*, *accessible content*, and the focus of attention. Note that the system's design makes a principled distinction between aspects of cognition that are and are not able to be introspected or verbally reported. Briefly, *uninspectable* processing occurs in the system's components whereas potentially *inspectable*, reportable representations are housed in its accessible content. Here, we briefly summarize the organizational details of ARCADIA reported by Bridewell and Bello (2015).

The representation of data and information in ARCADIA is multifarious and differs between its uninspectable and inspectable layers. Whereas components may freely use any representations and algorithms that facilitate their operations, effective communication among components necessitates an organizational system so that their results may be located and operated upon by others. Components share their results through interlingua elements that they place in accessible content. The current focus of attention is a privileged element in accessible content.

Information within ARCADIA is processed in the cognitive cycle illustrated by figure 1. On each cycle, the system broadcasts the focus of attention to all the components. To carry out their processing, those components may poll sensors or query accessible content. When finished, the components report interlingua elements to a temporary location where ARCADIA examines them, applies a focus selection strategy to identify the next focus of attention, and populates accessible content. Note that accessible content is flushed and replenished on each cycle, so unless some components act as system memory, information about past states is lost.

ARCADIA's design reflects knowledge about attention in humans, most importantly that it directs perceptual and cognitive resources. Much of the knowledge about attention is derived from errors associated with failures to attend. Consequently, we expect models in ARCADIA to fall prey to similar errors, including inattentional blindness. Further, inasmuch as those models are motivated by theories of attention, perceptual processing, and cognition, we claim that they provide much needed explanatory accounts for those errors. To that end, after providing an overview of ARCADIA's vision model, we introduce experiments on inattentional blindness taken from the literature and describe an extended model that accounts for key phenomena in those experiments.

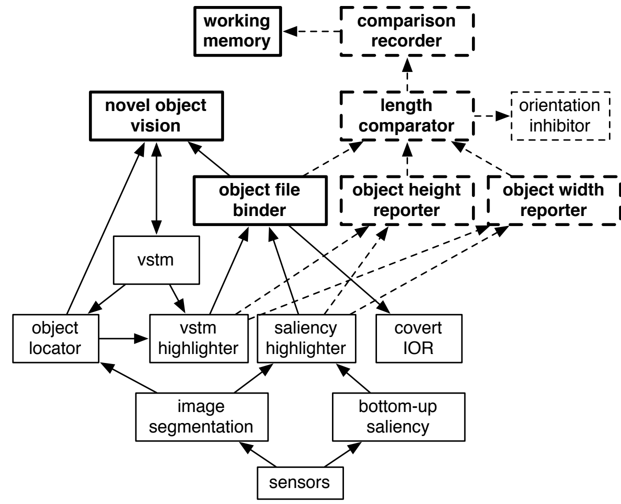

Figure 2: Component interactions in an ARCADIA model of visual processing. The boxes (except for "sensors") indicate the individual components. The arrows illustrate components that use the information output by others. All components operate over their internal state and accessible content, but those in bold font respond to the focus of attention. Dashed lines indicate parts of the model that are specific to the inattentional blindness task.

Vision in ARCADIA

Bridewell and Bello (2015) previously motivated and described a vision pipeline in ARCADIA, here we review that pipeline with a greater emphasis on the details of the components and how they work together to incrementally construct visual objects. Although components in ARCADIA are "always on" and operate in parallel, some components depend on the output of others. Referring to figure 2, the distinction between bottom-up and top-down interactions is roughly shown through the layers of components and the direction of the arrows. For instance, object locator, which tracks recently attended objects, receives bottom-up information from image segmentation and top-down guidance from visual short-term memory (vSTM).

ARCADIA's visual model is designed for object localization, identification, and tracking. To this end, the attentional strategy reflects priorities for constructing objects from visual data and updating their positions over time. ARCADIA differs from most cognitive architectures in that it operates over pixel-level, video data. Further, the system takes multiple cycles to construct object representations from low-level features (e.g., closed contours, color profiles, "retinal" location). Once constructed, the representations must be tagged as referring either to new object instances or to previously seen objects. Because objects take multiple cycles to encode, interruptions in the normal object-construction process may result in errors that are characteristic of human perception.

As shown in figure 2, vision in ARCADIA involves a variety of components. Space precludes discussing the computational details of each one, so in this paper, we highlight each component's role and the information it communicates to others. Beginning at the base of the figure, ARCADIA incorporates sensors. For the purposes of this paper, there is one simulated camera that operates over a video file. This camera reports rectangular video frames and roughly corresponds to the system's view of a computer screen. Since the modeled experiments assume that gaze rests at a central fixation point, only covert visual attention (i.e., without saccades) is used.

The sensor feeds into visual components that process their input without any top-down influence. The first of these implements the *bottom-up saliency* calculations described by Itti, Koch, and Niebur (1998) to determine where visual attention will shift in the absence of any other information. The second of the components carries out *image segmentation* and collects information about the color and location of closed-contour regions that might correspond to objects. The resulting regions are treated as proto-objects (Rensink, 2000) that may become full-fledged objects with attention.

The next set of components operates over covert visual attention. The *saliency highlighter* reads the interlingua elements created by bottom-up saliency and image segmentation to pick out a few (here, two) proto-objects that appear at points of high saliency (Xu & Chun, 2009). If such a proto-object exists, the component produces an interlingua element representing a candidate location for visual attention. Similarly, the *vSTM highlighter* requests visual attention at the location of any object currently stored in vSTM. Working in tandem with these components, *covert inhibition of return* (McDonald, Hickey, Green, & Whitman, 2009) discourages ARCADIA from attending to only one location so that it can assemble information about multiple objects within the scene.

To this point, the components have read from sensors or accessible content on each cycle, ignoring the information in the focus of attention. Moving upwards in figure 2, if the output of one of the highlighters becomes the focus of attention, the *object file binder* responds. This component takes the proto-object at the attended location and constructs an object-file representation for storage in vSTM.

Once an object-file exists, ARCADIA can determine whether it corresponds to a new object or one that it has already seen. Currently, this determination is made only for representations in vSTM. If the object file receives focus, *novel object vision* checks whether it is identical to elements tracked by vSTM. The establishment of object identity is a complex topic and depends on the task at hand and the properties of the objects themselves. ARCADIA can examine a variety of features to distinguish objects including size and color, but in this model, only location is used. Specifically, an object is treated as an update to a memory representation if its location overlaps the most recent position of the remembered object. Novel object vision generates either a statement of equality between the focal object and one in vSTM or an assertion of a new object.

The *vSTM* component searches accessible content for these statements of equality or newness and updates its memory stores accordingly. This component holds up to four object files at a time, updating them when possible, and replacing an arbitrarily selected old object file with any new one that receives focus. Because ARCADIA components encapsulate information, vSTM reports interlingua representations of its contents on each cycle. Consumers include novel object vision, which uses the elements for comparison purposes, vSTM highlighter, which requests covert orientation, and *object locator*. The design of object locator is based on *fingers of instantiation* (see, Scholl & Pylyshyn, 1999) and is not discussed here because the objects are static. Details are provided by Bello and colleagues (Bello, Bridewell, & Wasylyshyn, 2016).

This walkthrough of ARCADIA's model of visual processing describes the components and their informational exchange but has not discussed attentional strategies, which are task specific. With that caveat, we outline a limited attentional strategy for visual perception.

1. If an object file was just created, focus on it to enable comparison to and/or encoding in vSTM.
2. Else, if a highlighter directs covert visual attention to an uninhibited location, attend to the specified proto-object.
3. Else, attend to an arbitrarily selected interlingua element.

For static scenes, this strategy is insufficient for viewing more than two or three objects because saliency is the principle means for capturing visual attention. Introducing overt visual attention, a memory for previously viewed locations, or applying the strategy to dynamic scenes can reduce the tendency to stare at one or two objects. But, under the task constraints for the experiments we are modeling, even this basic strategy (with a couple task-related augmentations) works surprisingly well. The next section describes experiments on inattentional blindness in humans, after which we introduce a handful of new components and a corresponding attentional strategy to guide focus selection.

Experiments on Inattentional Blindness

Perhaps the most systematic examination of inattentional blindness was reported by Mack and Rock (1998) in their book-length treatment of the phenomenon. Over the course of several years, the authors collected thousands of data that supported the statement, "There is no explicit perception without attention." The studies were carried out across two laboratories, and although the methodology differed according to the hypothesis, the experimental protocol for stimulus presentation was generally consistent.

Most experiments consisted of nine trials per subject. Each trial followed roughly the same structure, although several variations were explored. For brevity, we describe the ex-

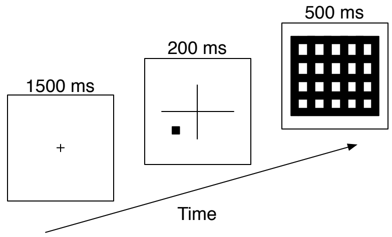

Figure 3: The time course of a single trial in the Mack and Rock (1998) studies.

perimental conditions for the specific study of interest. In this study, subjects were given the task of determining which of the two lines of a cross presented at fixation were the longest. On a standard trial the time course followed that shown in figure 3, a fixation point appeared for 1500ms, followed immediately by a stimulus presented at 200ms, and then a mask meant to overwrite iconic memory was shown for 500ms. After this presentation, the subjects were asked to report whether the horizontal or vertical line was longer.

During the first two trials, only the cross appeared. On the third trial, the cross was joined by a *critical stimulus* unrelated to the task. In this case, the critical stimulus was a black square that appeared in one of the four quadrants defined by the cross and within the rectangle implicitly defined by the length and the width of the cross. After this trial, the subjects were asked if they saw anything out of the ordinary. Three similar trials followed where the subject was asked to perform the original task while also reporting seeing any other objects. In the last set of three trials, the subjects were instructed only to report objects apart from the cross.

In this design, the critical stimulus was presented on the third, sixth, and ninth trial. During the first three trials, the subjects were considered to be naive to the presence of the critical stimulus. During the second three trials, the subjects were considered to have *divided attention*, and on the last three trials they were assumed to dedicate *full attention* to detecting the critical stimulus. Mack and Rock (1998) found that under various conditions roughly 25% of the people failed to see the critical stimulus on the third trial although almost 100% saw it during the ninth (full attention) trial.

The ability to account for perceptual differences across the various trials would be a notable success for ARCADIA's visual processing model. In particular, when the critical stimulus is present, the system should be able to create object files for the cross and the square, record those representations in vSTM, carry out the length comparison, and store the results in working memory. Because the task stimulus appears for only 200 ms, this gives the system few cycles in which to carry out this activity. Although we make no strong claims about the relationship between cycles within ARCADIA and physiological measures, we note that visual attention is linked to gamma band synchronization in the 40 to 80 Hz range (Fries, 2009). Taking the slow end of that range as the point where attended representations may shift, this would give the system 8 cycles to view and process the stimulus.

Modeling Inattentional Blindness

In the introduction, we pointed out four characteristics of attention that studies on inattentional blindness reveal. In this section, we report a model of attention that exhibits all four.

1. Cognitive tasks interfere with perception.
2. Attention is required for object representations to persist.
3. Object properties that do not persist cannot be reported.
4. Visual perceptions are processed without attention.

Furthermore, because the model receives video input that matches a rate of one frame per 25 ms, it accounts for these characteristics within a strict time limit. The task-relevant stimulus appears for only eight frames, giving ARCADIA eight cycles to orient to the visual objects, construct object files, carry out the comparison task, and memorize the result.

For this model, we emphasize those trials from the Mack and Rock experiments where the critical stimulus appears. Specifically, we will offer an explanation for why people (a) are not inattentionally blind in the full-attention condition, (b) are generally not inattentionally blind in the divided-attention condition, and (c) may or may not be inattentionally blind during the critical trial. We do not model the trials without the critical stimulus because we are not yet accounting for the process that lulls subjects into inattention.

Modeling the Mack and Rock experiments requires implementing the components with dashed borders in figure 2 and extending the attentional strategy accordingly. The changes include components for comparing object dimensions and a means for introducing a task-related inhibition of attention. In the rest of this section, we describe the new components, discuss the attentional strategy, and report the resulting behavior. We then address how the model exhibits the characteristics of inattentional blindness.

Additional Components

We developed five new components to carry out the task of comparing line lengths and remembering the result. Referring to figure 2, the *object height reporter* and *object width reporter* respond to the output of the highlighting components. In parallel with the construction of the object file representation, these components represent height and width in a comparable format. If the object file receives focus, then the *length comparator* carries out the straightforward comparison between width and height and reports its result. To ensure that the lengths and widths of other objects are not overtly compared (since that was not part of the task instructions), length comparator only operates on objects larger than the fixation mark. If a comparison was made and receives focus, then the *comparison recorder* requests that the result be memorized.

Working memory responds to attended requests for memorization, storing the associated representation for later use. Note that elements in working memory are not processed in this model, and the current implementation of working memory is little more than a storehouse. Elements in working memory are reported to accessible content on each cycle, as

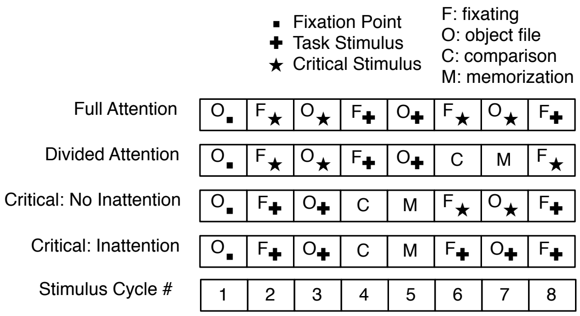

Figure 4: Focus traces for the four attentional conditions.

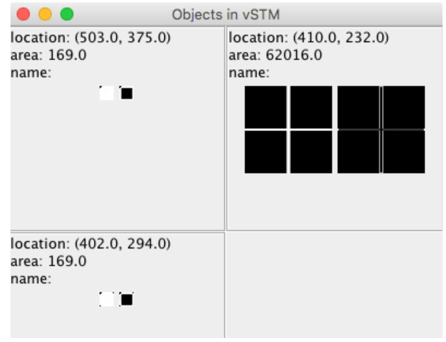

Figure 5: The resulting objects in vSTM for the divided attention trial.

with vSTM, but in the current model, no other components process them.

The final task-specific component is *orientation inhibitor*. This component plays a central role because it provides task-based limits to the parts of the visual field that receive attention. On each cycle, orientation inhibitor reports a request to inhibit non-central regions of the visual field. This reflects the expectations that the cross will appear at fixation and that nothing else is important to the comparison task. There are three settings to this component.

1. Off: no inhibition signal is sent. This is the assumed operative mode for divided attention and full attention trials.
2. Slow: an inhibition signal is sent on each cycle and the component takes several cycles to revoke inhibition after task completion. This is the assumed operative mode for the critical trial when there is inattentional blindness.
3. Fast: an inhibition signal is sent on each cycle and the component revokes inhibition immediately after task completion (i.e., after a comparison element receives focus). This is the assumed operative mode for the critical trial when there is no inattentional blindness.

Note that if inhibition is revoked quickly enough, the model cannot distinguish between the Fast and Off conditions for the critical trial when there is no inattentional blindness. However, this condition shows that even if task-based inhibition is active, there is time to perceive the critical stimulus.

Attentional Strategy

The attentional strategy for this task builds on the one provided earlier in the paper. Two new kinds of elements take priority over the other preferences in that strategy. They are (1) if there is an element that encodes a comparison result, make it the focus so that it can be memorized; and (2) if there is a request to memorize some information, focus on it in order to initate memorization. To review the rest of the strategy, if there are no comparisons or memorization requests, the strategy selects, in this order of preference, (3) a newly available object file, (4) a location where it should shift its covert attention, (5) an arbitrarily selected interlingua element. These changes allow the system to read in a video that represents a trial and output a comparison result and the objects that were perceived.

Results and General Discussion

There are four conditions of interest for a model of inattentional blindness: (1) the full attention trial without a task, (2) the divided attention trial with a comparison task and object detection, (3) the critical trial where the subject is not inattentionally blind, and (4) the critical trial where the subject is inattentionally blind. Figure 4 shows the results from running ARCADIA under these conditions. Each row in the figure depicts a separate condition, starting with the cycle where the task stimulus first appears. In all the trials, the system is in the middle of constructing an object for the fixation point, although it could just as well have been shifting its covert attention to a new location (the timing in this case would not alter the qualitative results). Simultaneously during cycle 1, the image segmentation component is producing elements that will be used by the highlighters.

In cycles 2 and 3, attention in the trials without task-based inhibition is drawn to the critical stimulus, whereas attention in the other trials selects the cross. In cycles 4 and 5 of the full-attention and divided-attention trials, ARCADIA processes the cross. The system's attention in the full-attention trial then oscillates between the critical stimulus and the cross due to covert inhibition of return. Attention in the divided-attention trial moves on in cycles 6 and 7 to compare the cross dimensions and memorize the result. This condition shows that without inhibition it is possible to carry out the task and perceive the entire visual scene within 200 ms. Figure 5 shows that after the divided attention trial, vSTM contains the fixation point, the cross, and the critical stimulus.

Turning to the critical trial conditions, task-based inhibition directs ARCADIA's attention to the central cross when it appears and completes the task. If inhibition of the peripheral regions is lifted in time, the system can then process the critical stimulus. If not, then the cross continues to receive attention. These results show that even if the system is directing its attention away from the critical stimulus, there is time to complete the task and encode all the visual objects. However, inhibition may take hundreds of milliseconds to revoke, blocking the encoding of the critical stimulus and, therefore, its explicit perception.

From the outset, we stated that a model of attention must exhibit four characteristics that arise in studies of inattentional blindness. First, the different focus traces for each condition show that cognitive tasks alter the time course of perception in ARCADIA. Second, without attention, objects fail to persist as attention is required for the construction of object files and for recording task results. Third, object properties that do not persist cannot be reported. In this case, the relationship between the cross's height and width is not remembered in the full attention trial and is subsequently not reportable. Fourth, although not all visual objects are explicitly perceived, they all undergo preattentive processing. Furthermore, this model provides an account for why otherwise salient objects may go unseen.

Concluding Remarks

To be sure, researchers working with cognitive architectures have addressed certain aspects of attention, perception, and cognition, but their intersection is rarely explored. On one side of a coin, the majority of computational work on vision and especially on object recognition is divorced from the psychological literature and interactions with cognition. On the other side of the coin, psychologically plausible approaches to perception in some of the better-known cognitive architectures such as ACT-R (Nyamsuren & Taatgen, 2013) and EPIC (Kieras, 2010) rely on symbolic encodings of objects, relations, and events. As a result, there seems to be an empty space between these two families of research approaches that an architecture like ARCADIA can fill: a space for systems that operate over real sensor data, but do so in psychologically plausible ways.

In this paper, we demonstrated that emphasizing task completion produces a sort of tunnel vision that blocks out otherwise salient stimuli. We suggest that this finding applies to any system that can control its focus of attention. Part of what it means to attend to one thing is to ignore everything else.

Acknowledgments

The authors would like to acknowledge generous support from the Office of Naval Research under grants N0001414WX20179 and N0001415WX01339. The views expressed in this paper are solely the authors and should not be taken to reflect any official policy or position of the United States Government or the Department of Defense.

References

Bello, P. F., Bridewell, W., & Wasylyshyn, C. (2016). Attentive and pre-attentive processes in multiple object tracking: a computational investigation. In *Proceedings of the thirty-eighth annual meeting of the cognitive science society.* Philadelphia, PA.

Borji, A., & Itti, L. (2013). State-of-the-art in visual attention modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *35*, 185–207.

Bridewell, W., & Bello, P. F. (2015). Incremental object perception in an attention-driven cognitive architecture. In *Proceedings of the thirty-seventh annual meeting of the cognitive science society* (p. 279-284). Pasadena, CA.

Drew, T., Vo, M. L.-H., & Wolfe, J. M. (2013). The invisible gorilla strikes again: Sustained inattentional blindness in expert observers. *Psychological Science*, *24*, 1848–1853.

Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annual Review of Neuroscience*, *32*, 209–224.

Hyman, I. E., 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.

Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *20*, 1254–1259.

Kieras, D. (2010). Modeling visual search of displays of many objects: the role of differential acuity and fixation memory. In *Proceedings of the tenth international conference on cognitive modeling* (pp. 127–132). Philadelphia, PA.

Mack, A. (2008). Inattentional blindness: Reply to commentaries. *Psyche*, *7*, 1–6.

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

McDonald, J. J., Hickey, C., Green, J. J., & Whitman, J. C. (2009). Inhibition of return in the covert deployment of attention: evidence from human electrophysiology. *Journal of Cognitive Neuroscience*, *21*, 725–733.

Nyamsuren, E., & Taatgen, N. A. (2013). Pre-attentive and attentive vision module. *Cognitive Systems Research*, *24*, 62–71.

Rensink, R. A. (2000). Seeing, sensing, and scrutinizing. *Vision Research*, *40*, 1469–1487.

Scholl, B. J., & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: clues to visual objecthood. *Cognitive Psychology*, *38*, 259–290.

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

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*, 722–727.

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting memories: cognition of brief visual stimuli.* (pp. 71–94).

Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in Cognitive Sciences*, *13*, 167–174.

Yantis, S., & Egeth, H. (1999). On the distinction between visual salience and stimulus-driven attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *25*, 661–676.