Characterizing a snapshot of perceptual experience

Michael A. Cohen^{1,2}, Caroline Ostrand², Nicole Frontero², Phuong-Nghi Pham³

¹ McGovern Institute for Brain Research, Department of Brain and Cognitive Sciences,
Massachusetts Institute of Technology

² Department of Psychology, Amherst College

³ Department of Psychology, Ryerson University

Abstract

What can we perceive in a single glance of the visual world? While this question appears rather simple, answering it has been remarkably difficult and controversial. Traditionally, researchers have tried to infer the nature of perceptual experience by examining how many objects and what types of objects are not fully encoded within a scene (e.g., failing to notice a bowl disappearing/changing). Here, we took a different approach and asked how much we could alter an entire scene before observers noticed those global alterations. Surprisingly, we found that observers could fixate on a scene for hundreds of milliseconds yet routinely fail to notice drastic changes to that scene (e.g., scrambling the periphery so no object can be identified, putting the center of one scene on the background of another scene, etc.). In addition, we also found that as observers allocate more attention to their periphery, their ability to notice these changes to a scene increases. Together, these results show that although a single snapshot of perceptual experience can be remarkably impoverished, it is also not a fixed constant and is likely to be continuously changing from moment to moment depending on attention.

Introduction

What is the resolution of a snapshot of perceptual awareness? What can observers perceive “in the blink of an eye?” To answer these questions, researchers have often used paradigms like change blindness and inattention blindness to determine how much information observers fail to notice in a scene (Rensink et al., 1997; Mack & Rock, 1997; Simons & Chabris, 1999). In a typical experiment, individual items within a scene change or appear in some unexpected manner (i.e., a person’s sweater will change color, a chair will disappear, a man in a gorilla suit will appear, etc.) and the frequency with which observers notice those items are measured. Because so many observers routinely fail to notice these occurrences, many researchers agree that observers are aware of surprisingly little of the world around them and perception is rather sparse (Dennett, 1991; Baars, 2002; Simons et al., 2005; Cohen et al., 2012; Dehaene & Changeux, 2016), though of course, this view is extremely controversial (Lamme 2003; Block, 2007; Koch & Tsuchiya, 2007; Haun et al., 2017).

While these results are striking and informative, it is widely believed that there are firm limits as to how much information can go unnoticed when looking upon a scene. For example, in a typical change blindness or inattention blindness experiment, the item that changes or appears is usually relatively simple and its change or appearance does not fundamentally alter the meaning (or “gist”) of the scene. Indeed, it has been routinely claimed that if a change does affect the gist of the scene, that change will be easily noticed (Simons & Levin, 1997; Mack & Rock, 1997; Rensink, 2000; Koch & Tsuchiya, 2007; Sampanes et al., 2008; van Boxtel et al., 2010; but see Cohen et al., 2011 and Mack & Clarke, 2012). Intuitively, this idea makes quite a bit of sense: while it is easy to imagine a person failing to notice the keys on the dining room table, it is difficult to imagine how they could fail to notice that their dining room has been changed into a glacier (Lamme 2003; Block, 2007).

The idea that perception of the scene as a whole is robust, even if individual items in that scene are often not encoded, is bolstered by numerous results that have shown how several dimensions of a scene can be extracted from just a single fixation. For example, with just a few hundred milliseconds of exposure, observers are able to perceive the gist of a scene, identify a few prominent objects within that scene, and describe a variety of properties such as the navigability, depth, and temperature of a scene (Potter, 1976; Intraub, 1981; Fei-Fei et al., 2007; Castelhamo & Henderson, 2008; Greene & Oliva, 2009; Walther et al., 2009; Torralbo et al., 2013). Together, these results raise a natural question: if observers routinely fail to perceive individual objects within a scene, but appear are typically aware of general properties of the scene around them, what exactly are the limits of perceptual awareness in natural scenes? How much can observers truly encode in a single “snapshot” of perceptual experience?

To answer these questions, we took a novel approach. Rather than simply alter individual objects within a scene and measure if observers noticed those alterations, we asked how much of an *entire scene* we could alter before observers noticed. The logic of this approach is that if observers do not realize we made certain alterations to the world, those alterations were not a part of their perceptual experience. However, if observers do notice those alterations, we have removed important elements of their subjective experience of the world. By systematically examining what we can and cannot change in a scene before observers notice, we can quantify and visualize a snapshot of perceptual experience.

The approach used in this set of studies was directly inspired by a series of pioneering experiments done by Rayner and colleagues that focused on the size of the perceptual span in reading (McConkie & Rayner, 1975; Rayner, 1998; Rayner, 2014). In these studies, the researchers used a gaze-contingent technique such that only the parts of the text that an observer is fixating on is presented normally, while the text just beyond the fovea is degraded (i.e., replaced by strings of X's, out of focus, etc.). In this case, rather than move our window through the visual field, we simply locked it in place in the middle of the display and kept our observers' eyes fixated on that location. Specifically, we used the notion of preserving the foveal region of an image to determine how much of a natural we could degrade beyond the fovea without observers noticing. While the gaze-contingent technique has since been widely used to separate foveal and peripheral processing in a variety of situations such as visual search (Pomplun et al., 2001; Reingold et al., 2001) and scene perception (Loschky & McConkie, 2002; Miell et al., 2010; Võ & Henderson, 2011), the results reported here are the first to use this approach to directly probe the limits and bandwidth of perceptual experience.

Across three experiments, we used a modified inattention blindness paradigm to determine how often observers will notice different types of alterations to images of natural scenes while performing a variety of perceptual tasks. In each experiment, we employed the following basic procedure. Observers performed a primary task for a few trials that involved looking at large images of scenes (i.e., 26° of visual angle) and doing a basic object detection or scene classification task. Then, on the last trial, without observers knowing anything different would occur, we presented a critical stimulus that had been altered in the periphery in one of several ways. For example, in some cases, the periphery was so scrambled that no object could be detected or identified. In another case, the

center part of one image (e.g., a table in a dining room) was surrounded in the periphery by an inconsistent image (e.g., a glacier). In all of these cases, the only part of the images that were unaltered corresponded to the fovea/parafovea, which is a radius of approximately 4° of visual angle (Polyak, 1941; Swanson & Fish, 1995; Millodot, 2014). To put that number in perspective, in everyday life, the fovea/parafovea comprises less than 1% of the visual field¹. Thus, our alterations distorted an overwhelming majority of the stimuli. How often would such drastic alterations be noticed?

Surprisingly, we found that a large number of observers could fixate upon an image for several hundred milliseconds without noticing a wide variety of radical alterations made to the stimulus in the periphery. For example, a large portion of observers were completely oblivious that a scene they had just gazed upon was so degraded in the periphery that no single object could be recognized or identified. In addition, a majority of observers were totally unaware of the fact they had looked directly at a scene in which the fovea/parafovea region came from one image (i.e., a dining room), while the periphery of that image came from an entirely different image (i.e., a glacier). Thus, contrary to prior claims that observers are always aware of the gist of the scene and can rapidly extract a variety of dimensions about a scene, we find several situations where this is simply not the case. Indeed, in our most drastic situations, observers are not simply unaware of individual items within a scene; they are unaware of the fundamental changes we have made to the global properties of the entire scene.

However, it should be stressed that even though observers were not aware of these extreme alterations, they were still able to process enough of the stimulus to pick it out of a lineup (i.e., “I saw that woman right there, not any of those other people,” Figure 1). In other words, observers could fixate on a scene, process enough of that scene to identify the item they had just fixated upon, but still completely fail to realize that the vast majority of the scene was completely scrambled or replaced by another scene altogether.

Finally, it should also be stressed that the resolution of perception we are describing here does not hold in all situations at all time. In another experiment, we found that the size of these effects varied significantly as a function of the way participants allocated their attention across the scene. In our first two experiments, participants performed a simple object detection task at fixation (i.e., Is a human face present or absent?). In the third experiment, participants performed a task that drew their attention more towards the periphery (i.e., Is the scene of an indoor or outdoor image?). In this case, we found that when attending more towards the periphery, observers were more likely to notice the critical stimuli.

Taken together, these results highlight how a single snapshot of perceptual experience may be far more impoverished than is widely believed. While it has previously been shown that individual objects or items can be changed/replaced without observers noticing, these are the first results showing how an entire scene can be radically altered without them noticing. Moreover, the fact that these effects varied as a function of attention also highlights how the resolution of perceptual experience is not a static constant. These results have serious implications for many theories of visual cognition and perceptual awareness.

Experiment 1: Scrambled periphery

Here, we asked how much of the periphery we could scramble before observers would naturally notice. In this case, we systematically altered the extent to which the periphery of an image was scrambled while preserving certain basic natural scene statistics. At its most extreme level, images were so severely scrambled that individual objects in the periphery were completely unidentifiable. How often will observers notice such drastic alterations?

Participants

A total of 40 participants were run in each of the three experimental conditions, resulting in a total of 120 participants. All participants were recruited from the Amherst College and Massachusetts Institute of Technology communities. All participants gave informed consent and were compensated with either \$5 or course credit. All experimental procedures were approved by the Committee on the Use of Human Subjects in Research under the institutional review board of Amherst College.

Procedure

Participants performed a simple face detection task at fixation. On each trial, they were shown between 7-30 images of natural scenes and reported whether the last image they saw in the stream contained a face in the middle of it (Figure 1). The number of images shown on every trial was randomly sampled from between 7 and 30. Each image was shown for 288ms with a 100ms gap in between. We chose a presentation rate of 288ms/item because an analysis of several eye-tracking studies revealed that this is approximate duration of single a fixation period in naturalistic viewing conditions (Rayner, 1998; Pelz & Canosa, 2001; Hayhoe et al., 2003; Henderson, 2003; Unema et al., 2005; Rayner et al., 2008; Castelhana & Heaven, 2011; Nuthmann, 2017; see Appendix). Regardless of whether or not the trial had a face at the end of it, there were anywhere between 2 and 5 images that had faces in the middle throughout the stream. Participants were explicitly told that this would happen and that they were only to say whether or not the *last* image had a human face in the middle. All images were randomly selected from trial to trial from the different stimulus folders (i.e., “face targets” and “no face distractors”). No image was ever seen twice within a single trial, although images could be seen multiple times across trials. For the first 10 trials, half of the trials had a face target present at the end and half did not. At the end of each trial, a screen appeared that prompted the participant to say whether or not the last image had a human face in the middle. Participants pressed one key to report that a target face was present and another key to report that no target face was present.

Figure 1. Visualization of the trial procedures for Experiments 1 and 2. Participants performed 10 trials where they simply said if the last image in the stream did or did not contain a human face in the middle. Then, on trial 11, an unexpected critical stimulus was presented at the end of the trial and participants were immediately probed to determine whether or not they noticed the critical stimulus. Finally, participants were shown four examples of different critical stimuli and asked to choose which image they saw at the end of the last trial (in most cases, participants reported picking the image based on the preserved/middle part of the image).

On the critical trials, the selection procedure for the number of stimuli used in the entire trial (i.e., 7-30) and the number of face target stimuli used throughout the trial (i.e., 2-5) was identical to all previous trials. The main difference was primarily that the very last image in the stream was a critical stimulus. Within every experimental condition of both experiments, half of the participants were shown a critical stimulus with a human face in the middle and the other half were shown a critical stimulus that did not have a human face in the middle. As soon as the critical stimulus disappeared, the display went blank. No text was presented to prompt the participants to give any response. Instead, the moment the trial ended, an experimenter would immediately begin asking a series of questions to the participant. Those questions were:

- 1) "Did you notice anything strange or different about that last trial?"
- 2) "If I were to tell you there was something different about that last trial, could you tell me what it was?"
- 3) "If I were to tell you there was something different about the very last image you saw on that last trial, could you tell me what it was?"

If participants could not correctly identify the critical alteration after any of those questions, the trial ended with the four alternative forced choice task where they were asked to identify the specific image they saw. In this case, observers would be shown four stimuli that had been altered in the same way (i.e., four images that were medium scramble, four images that had an abstract periphery, etc.). In addition, all four images would either all contain a face target in the middle or none of them would contain a face target in the middle depending on the specific stimulus participants saw on the last trial. Upon being shown those four images, participants were told what had happened on the last trial and were asked to confirm whether or not they had noticed. Only those participants who claimed to not notice the critical manipulation when seeing the four images were classified as being completely inattentionally blind and were reported in the main figures. Finally, participants were asked to select the image they saw on the last trial

by pressing one of four different keys on the keyboard to indicate which image they saw on the last trial (i.e., “I saw the image with this particular person, not those images with those other people.”). Showing four examples like this served a few purposes. First, it allowed us to show participants the alterations we made to the images and explicitly verify that they did not notice those modifications. Second, it is possible that many of our participants were unaware of the critical manipulations because they did not notice the last image at all for one reason or another (e.g., sneezing, blinking, or a phone ringing at an inopportune time). Thus, we wanted to determine how often participants did not notice the alterations but did process and remember the critical stimulus.

Stimuli

All stimuli were shown on a 27-inch Apple iMac computer with a refresh rate of 60Hz and 2560 x 1440 pixels. A chin and forehead rest was used to minimize head movements and maintain a constant viewing distance of 57cm. The experiments were created and controlled using MATLAB and the Psychophysics toolbox (Brainard, 1997; Pelli, 1997). The background color of the screen was gray, with a luminance of 68.5 cd/m². Each individual stimulus subtended 26° of visual angle. For critical stimuli, the center part of the image was preserved by a flat top Gaussian that subtended 4° of visual angle (radius). The flat top portion of the Gaussian subtended 2.5° of visual angle (radius), which approximately corresponds to the fovea and was entirely unaltered. Meanwhile, the remaining 1.5° of visual angle were the slopes of the Gaussian served as a gradual fade from the untouched center of the image to the altered periphery of the image. This sloping/faded region comprised the remaining portion of the preserved region and would land approximately on the parafoveal part of the retina. All initial stimuli were acquired through online resources (i.e., Google) and personal photograph collections of the authors and were then converted to greyscale. See Supplementary Figures 1-4 for multiple examples of stimulus generation or procedures and example stimuli.

Face task stimuli: For the face detection task, there were 30 face target images and 30 distractor images. Each face target had a human face that was unambiguously in the middle of the image, viewed either head on or from the side. Individual faces came from a variety of genders, ages, and ethnicities. Each no-face distractor was defined by not having a human face in the middle. In other words, there could be animal faces in the middle and humans in the periphery, but without a human face in the middle, those images were considered distractors. Before the experiment began, observers were shown multiple examples of both target and distractor images to understand what differentiated targets and distractors. For both face and no-face stimuli, the images comprised a wide variety of settings and locations (e.g., a man in a kayak, a construction worker, a chef in a kitchen, a dog on a dog bed, an empty office, an empty beach, etc.).

Figure 2. Representative examples of both target and distractor stimuli used in Experiments 1 and 2.

Critical stimuli: All critical stimuli were generated using texture synthesis algorithms that preserve a variety of low-level image statistics within pre-specified pooling windows (Rosenholtz, 2016). In this case, we used three different pooling window sizes for our three experimental conditions: 1) *Small scramble*: Stimuli in this condition were generated using the Freeman and Simoncelli (2011) algorithm with a scaling parameter of 0.5, which corresponds to the ratio of the diameter of a pooling window relative to the eccentricity of that pooling window in the periphery (Figure 3). 2) *Medium scramble*: Stimuli in this condition were also generated using the Freeman and Simoncelli algorithm, except this time they were created with a scaling parameter of 3.0, resulting in images that appeared far more scrambled than in the small scramble condition. 3) *Large scramble*: Stimuli in this condition were generated using the Portilla and Simoncelli algorithm, in which one large pooling window spans the entire image (2000). Stimuli from the large scramble condition are so thoroughly scrambled that no individual object can be identified. Finally, in each of the three conditions, after generating the stimuli with these different algorithms, we also replaced the central parts of the scrambled images with a preserved section of the original image. This preserved section subtended a circular aperture with a radius of 4° of visual angle, which roughly accounts for the entire foveal/parafoveal regions of the retina (Polyak, 1941; Swanson & Fish, 1995; Millodot, 2014; see Supplementary Figure 1). Thus, for each critical stimulus, the region of the image that landed on the fovea/parafovea was normal, while the information on the periphery was scrambled to varying degrees (for more examples of critical stimuli, see Supplementary Figures 2-4).

Figure 3. Representative examples of the different types of critical stimuli used on the critical trial of Experiment 1. Note that each image subtended 26° of visual angle in the experiment, with only the center 4° of visual angle (radius) being preserved.

Results

What percentage of participants will be unaware of the scrambling done to these images when they are attending at fixation? In the small scramble condition, 100% of participants (40/40) were unaware of the stimulus alterations (Figure 4, see Supplementary Materials for more detailed analyses of the data). This result is unsurprising since the scrambling is small relative to that in the other conditions, and thus is not always immediately apparent. However, this amount of scrambling is indeed detectable when observers attend directly to the periphery (Wallis et al., 2019). Somewhat surprising, however, is the fact that 50% of participants (20/40) were unaware of the stimulus alterations in the medium scramble condition. In fact, an even more surprising result is the

fact that in the large scramble condition, where the periphery of the stimuli are scrambled so much that no individual objects can be identified, 48% of participants (19/40) were completely unaware of the stimulus alterations. In this case, there was a significant difference in noticing rates between the small and medium scramble conditions ($\chi^2(1)=24.07, P<0.001$), as well as between the small and large scramble conditions ($\chi^2(1)=25.83, P<0.001$). Meanwhile, there was no significant difference between the medium and large scramble conditions ($\chi^2(1)=0, P=1$). Finally, performance on the primary face task was nearly perfect for each of the three experimental conditions (Small scramble: 97.3% correct; Medium scramble: 98.5% correct; Large scramble: 97% correct) with no significant difference in performance between any of the three conditions ($t(39)>1.22, P>0.22$ in all cases).

Figure 4. Inattention blindness rates for each condition in Experiment 1. The percentage of participants who failed to notice the critical stimulus is plotted on the vertical axis. Each bar corresponds to a different experimental condition. *** $P<0.001$.

As mentioned above, the final step in our experimental procedure was showing participants four examples of different critical stimuli in order to: 1) verify that they did not notice the stimulus alterations and 2) ask if they could identify the last image even if they were unaware of the critical manipulations. In this case, the vast majority of participants were able to identify the last image they saw in each experimental condition (small: 36/40 participants; medium: 34/40 participants; large: 36/40 participants). Even though these rates are quite high, it is still possible that they underestimate the percentage of participants who were aware of the last image. Since a good amount of time would sometimes elapse between the presentation of the critical stimulus and the presentation of the four options, it is possible that some participants were aware of the critical stimulus but simply confused it with another stimulus after the questioning period. Regardless, our next analysis focused on determining the inattention blindness rates when only examining the

data from participants who could identify the last image. In this case, we found that the inattentional blindness rates were virtually identical to when all participants' data was analyzed (small scramble: 100%, 36/36 participants; medium scramble: 43%, 14/34 participants; large scramble: 50%, 18/36 participants). In other words, in each condition, even though over 40% of participants processed enough of the critical stimulus to pick it out of a lineup of four options, they were completely unaware of the significant alterations that had been made to the stimulus in the periphery.

Experiment 2: Inconsistent, abstract, and no periphery

Perhaps the most interesting result in Experiment 1 is the fact that so many observers failed to notice the alterations to the stimuli in the large scramble condition. If several of our observers are unaware of those alterations, what does it take to get them to routinely notice stimulus alterations? In Experiment 2, we tried to answer this question by repeating the procedures of Experiment 1 with a new set of critical stimuli that altered the periphery in a variety of ways.

Participants

A total of 40 participants, none of whom participated in Experiment 1, were run in each of the three experimental conditions, resulting in a total of 120 participants. All participants were recruited from the Amherst College and Massachusetts Institute of Technology communities. All participants gave informed consent and were compensated with either \$5 or course credit. All experimental procedures were approved by the Committee on the Use of Human Subjects in Research under the institutional review board of Amherst College.

Procedure

The only procedural difference between Experiments 1 and 2 was with regard to the critical stimuli. Besides changing the critical stimuli, the procedures between the two experiments were identical.

Stimuli

There were three types of critical stimuli in this experiment: 1) *Inconsistent periphery*: The foveal/parafoveal part of one image is placed on top of a semantically inconsistent peripheral image (i.e., a part of a coffee shop placed on top of a desert scene, etc., Figure 5). 2) *Abstract periphery*: The foveal/parafoveal part of one image is placed on top of high-contrast abstract shapes and patterns. 3) *No periphery*: The foveal/parafoveal part of one image is placed on top of a uniform grey background (see Supplementary Figures 5-8 for a schematic on stimulus creation and more examples of the critical stimuli).

Figure 5. Representative examples of the different types of critical stimuli used on the critical trial of Experiment 2. Note that each image subtended 26 degrees of visual angle in the experiment, with only the center 4° of visual angle being preserved.

Results

How often will participants be aware of these critical stimuli when performing the same face detection task as in Experiment 1? The results from Experiment 2 are plotted in Figure 6. In the inconsistent periphery condition, 73% of participants (29/40) were unaware of the stimulus alterations on the critical trial. Put more simply, the vast majority of participants were unaware of the fact that they had seen stimuli comprised of parts from two entirely different scenes. It was not until the abstract periphery and no periphery conditions that the overwhelming majority of participants noticed the stimulus alterations (abstract periphery: 18%, 7/40; no periphery: 13%, 5/40). In this experiment, there was a significant difference in noticing rates between the inconsistent and abstract periphery conditions ($\chi^2(1)=22.27, P<0.001$), as well as between the inconsistent and no periphery

conditions ($\chi^2(1)=27.06, P<0.001$). However, there was no significant difference between the abstract and no periphery conditions ($\chi^2(1)=0.10, P=0.75$). As in Experiment 1, the majority of participants were able to identify the last image they saw in each experimental condition (inconsistent periphery: 35/40 participants; abstract periphery: 37/40 participants; no periphery: 37/40 participants). Furthermore, we again found that the inattentional blindness rates were virtually identical when only participants who could identify the critical image were analyzed (inconsistent periphery: 71%, 25/35 participants; abstract periphery: 16%, 6/37 participants; no periphery: 11%, 4/37 participants). As was the case in Experiment 1, performance on the primary face task was nearly perfect for each of the three experimental conditions (Inconsistent periphery: 98.5% correct; Abstract periphery: 96.8% correct; No periphery: 98% correct) with no significant difference in performance between any of the three conditions ($t(39)=1.30, P>0.02$ in all cases).

Figure 6. Inattentional blindness rates for each condition in Experiment 2. The percentage of participants who failed to notice the critical stimulus is plotted on the vertical axis. Each bar corresponds to a different experimental condition. *** $P<0.001$.

Experiment 3

Taken together, the results from Experiments 1 and 2 suggest that a single snapshot of perceptual experience is remarkably impoverished when observers are attending at fixation. Of course, in everyday life, observers can deploy their attention in a variety of ways. For example, the scale of attention will be radically different when searching for an individual in a crowded room compared to when staring at an open landscape and trying to take in its distributed layout. For the first two experiments, we chose a simple task that required participants to attend at fixation since it is generally the case that people attend to whatever they are fixating on and wherever they will fixate next (Posner, 1980; Deubel & Schneider, 1996; Hoffman & Subramaniam, 1995; Kowler, Anderson, Doshier, & Blaser, 1995). However, since attention can easily be directed more towards the periphery, we

sought to determine how the scale of attention affects a snapshot of perception. Thus, in Experiment 3, we replaced the face detection task with an indoor/outdoor task that was designed to naturally get participants to attend more to the periphery. In this case, rather than say whether the last image had a human face in the middle, observers simply said whether the last image was of an indoor scene or an outdoor scene. The logic behind this experiment is that by asking participants to make a judgment about the entire scene, rather than just about an object at fixation, their attention would naturally be distributed across the scene in a way that is somewhat ecologically valid. In addition, by distributing attention more towards the periphery, we are also bolstering the processing of peripheral vision, which has been shown to be more important for gist perception in a single fixation (Larson & Loschky, 2009). Thus, we predicted that the rates of inattention blindness would decrease in Experiment 3 relative to Experiments 1 and 2.

Participants

40 participants were run on the indoor/outdoor task. All participants were recruited from the Amherst College and Massachusetts Institute of Technology communities and none of whom participated in either Experiments 1 or 2. All participants gave informed consent and were compensated with either \$5 or course credit. All experimental procedures were approved by the Committee on the Use of Human Subjects in Research under the institutional review board of Amherst College.

Procedure

In Experiment 3, the procedures were identical to Experiments 1 and 2 with a few mild exceptions. First, an eye-tracker (i.e., a Gazepoint GP3 eye tracker) was used to ensure participants maintained fixation throughout the trial. Given the task demands of Experiments 1 and 2, it seems safe to assume participants maintained fixation throughout. In fact, if participants did not, that would only *decrease* our inattention blindness rates since participants are far more likely to notice critical stimuli if they are foveating on parts of the image we did not preserve. For Experiment 3, meanwhile, participants performed a 5-point calibration sequence at the beginning of the experiment. Once a trial began, if participants broke fixation (i.e., more than 2.5° of visual angle away from fixation), the trial would start over. This was fairly infrequent and occurred on less than 4% of trials.

For the first 10 trials of Experiment 3, the primary difference between this experiment and the previous two was simply that participants reported whether the last image was of an indoor or outdoor scene (Figure 7a). Again, before the experiment began, participants were shown multiple examples of indoor and outdoor stimuli so they could familiarize themselves with the task. In addition, for the last image of the trial, half of the trials had an indoor scene at the end and half had an outdoor scene at the end. Beyond those differences, the remaining procedures of Experiment 3 were identical to Experiments 1 and 2 up to and including the procedures of the critical trial.

a) Example indoor stimuli

Example outdoor stimuli

b) Face and scene tasks (different participants)

Figure 7. a) Representative examples of both indoor and outdoor stimuli used in Experiments 3. b) Visualization of the trial procedures for the face and scene tasks. Note that two groups of participants performed the face and scene tasks respectively.

Stimuli

For this experiment, we used two different types of stimuli. Specifically, we used stimuli from the medium scramble condition (Experiment 1) and the inconsistent periphery condition (Experiment 2) as our critical stimuli. For the indoor/outdoor task, the same stimuli were used as in the face task and the procedures used to select the stimuli from trial to trial were identical as the face task. Thus, the stimuli did not differ between the two tasks; only the tasks being performed on those stimulus streams were different. The lone exception to this rule came in selecting the stimuli for the last image in a stream. In certain cases, we decided that some of the images were somewhat ambiguous in terms of whether or not they were indoors or outdoors (e.g., an image of a woman working in a green house). To avoid any confusion, those ambiguous images were never shown as the last image in the stream (although they could be shown during a trial in any position that was not the last image in a stream). However, on the critical trial, no stimuli were excluded from consideration. The same exact stimulus batch and the same selection procedures used in the face task were used in the indoor/outdoor task.

Results

Overall, we found that the scale of attention had a dramatic effect on the resolution of perception. As a reminder, in Experiment 1, we found that 50% of participants (20/40) did not notice the medium scramble stimuli when performing the face task. By comparison, in Experiment 3, the inattentional blindness rates were significantly lower with only 15% of participants (6/40) failing to notice the critical stimuli when performing the scene task ($\chi^2(1)=9.63, P<0.01$; Figure 7b). When only examining the data from participants who could identify the last image, we found that the inattentional blindness rates dropped to 8% (3/37 participants). Meanwhile, a similar pattern of results was found when using the inconsistent periphery stimuli. Again as a reminder, in Experiment 2, we found that 73% of participants

(29/40) did not notice the critical stimuli when performing the face task. By comparison, in Experiment 3, the inattentional blindness rates were significantly lower with 45% of participants (18/40) failing to notice the critical stimuli when performing the scene task ($\chi^2(1)=4.22, P<0.05$; Figure 7c). When only examining the data from participants who could identify the last image, we found that the inattentional blindness rates dropped to 38% (13/34 participants). These results highlight how the resolution of perceptual experience is not a fixed constant, but is an ever-changing concept that varies as a function of attention.

Figure 8. a) Representative examples of the critical stimuli used on the critical trial of Experiment 3 with the medium scramble stimuli (and used in the medium scramble condition of Experiment 1), and inattentional blindness rates for the face and scene tasks. b) Representative examples of the critical stimuli used on the critical trial of Experiment 3 with the inconsistent periphery stimuli (and used in the inconsistent periphery condition of Experiment 3), and inattentional blindness rates for the face and scene tasks. In both cases, the percentage of participants who failed to notice the critical stimulus is plotted on the vertical axis. Each bar corresponds to a different experimental condition. ** $P<0.01$, * $P<0.05$.

When comparing inattentional blindness rates between primary tasks (i.e., face vs. scene tasks), it is important to also compare performance on those very primary tasks since differences in primary task performance can lead to differences in inattentional blindness rates. However, in this case we found no such evidence of a difference in primary task performance for the medium scramble condition (Face task performance – 98.5%; Scene task performance – 96.8%; $t(39)=0.39, P=0.70$) or the inconsistent periphery condition (Face task performance – 98.5%; Scene task performance – 97.5%; $t(39)=0.81, P=0.42$). Of course, these results should be interpreted with some caution since there is likely a ceiling effect in these cases with performance being above 96.8% in all cases.

General Discussion

When discussing the limits of visual cognition, it has been often claimed that even though observers may fail to encode individual objects, they will effectively always process the general scene around them (Simons & Levin, 1997; Mack & Rock, 1997; Rensink, 2000; Koch & Tsuchiya, 2007; Sampanes et al., 2008; van Boxtel et al., 2010; but see Cohen et al., 2011 and Mack & Clarke, 2012). This idea is consistent with the fact that

observers can not only recognize the gist of a scene very briefly (Potter, 1976; Intraub, 1981; Fei-Fei et al., 2007), but can also extract a variety of other higher-level dimensions about a scene as well (Tatler et al., 2003; Castelhamo & Henderson, 2008; Greene & Oliva, 2009; Walther et al., 2009; Torralbo et al., 2013). Across multiple experiments, however, we showed that there are many cases in which observers can fixate upon a scene for several hundred milliseconds, process and remember the items they were fixating upon, and yet still fail to notice shockingly drastic alterations that alter the entire meaning and interpretation of the scene. These findings suggest that awareness of the contents of peripheral vision in a single snapshot of perceptual experience can be surprisingly impoverished. Indeed, we found that a large portion of observers failed to realize that their periphery was so degraded that no object could be identified or that it was comprised of an entirely different scene. Of course, it should be emphasized that it is not always the case that observers will routinely fail to notice such drastic alterations. In Experiment 3, we found that significantly more observers noticed the scrambled or inconsistent periphery when their attention was deployed away from fixation in order to perform the indoor/outdoor task. Together, this collection of findings highlights how the resolution of perception can vary tremendously based on the deployment of attention. This is a crucial idea for researchers to keep in mind when trying to understand the richness of perceptual experience since it appears that there cannot be one definitive answer to the question, how rich is perceptual experience?

When considering these results, one prominent question naturally emerges: How can these findings be reconciled with personal introspection, which strongly suggests we see far more than these findings imply? We believe there are several possible answers to this question. For starters, in our experiments, observers are only presented with a single presentation of images that approximates the length of an average fixation period (288ms/item). In everyday life, however, observers do not see a single snapshot of a scene in isolation and then another snapshot of an entirely new scene. Instead, observers are usually embedded in an environment that they can look upon and explore for seconds, minutes, or even hours. One intriguing possibility is that observers build up a detailed representation of the scene around them over time, and the knowledge/memory of that representation provides the foundation for an apparently rich perceptual experience (Hollingworth, 2009). Alternatively, some researchers would argue that these results are consistent with the notion of “subjective inflation” (Knotts et al., 2019). Under this view, observers systematically overestimate the richness of their perceptual experience when attention is primarily focused elsewhere. Support for this idea comes from several studies showing that when objective performance is matched between two experimental conditions (e.g., cued vs. uncued stimuli), participants are more likely to mistakenly claim an item is present, when in reality it is not, whenever attention is reduced (i.e., in the uncued condition; Rahnev et al., 2011; Solovey et al., 2015). In the context of our results, observers might overestimate the richness of their peripheral experience and be unaware of the degree to which an image can be altered without their realizing it when attention is focused on the fovea during the face detection task. Of course, reconciling the results reported here with the subjective feeling of a richer perceptual experience is not a trivial endeavor, especially given the longstanding controversy surrounding this issue (Rensink et al., 1997; Mack & Rock, 1997; Simons et al., 2005; Koch & Tsuchiya, 2007; Lamme, 2010; Block, 2011; Kouider et al., 2011; Cohen et al., 2012; 2016; Dehaene & Changeux, 2016; Boly et al., 2017).

Fortunately, the results reported here raise a variety of questions that are far more tractable and could easily be addressed by future research. For example, earlier we discussed how one potential shortcoming of focusing on individual snapshots of perception is that it discounts the possible role of built up representations over time. One approach to address this issue would be to compare inattentional blindness rates with familiar vs. unfamiliar scenes. It is possible that if observers are already familiar with the critical stimulus, which was not the case in the present experiments, their pre-existing knowledge of that scene will enable them to spot alterations more easily. However, it should be noted that an alternative pattern of results could also emerge. If observers are familiar with a scene, their visual system may be able to “fill-in” certain details from memory and the inattentional blindness rates could actually increase (Komatsu, 2006). Another factor not addressed in this study, but which could be easily incorporated into future experiments, focuses on the richness of color perception. Another long running debate focuses on how much color observers perceive in their periphery (Tyler, 2014; Haun et al., 2017). The paradigm described here could be adapted to contribute to this controversy by making all of the images in color, rather than in black and white, and making the critical stimulus black and white in the periphery, and colored in the fovea. The extent to which observers notice such color alterations could shed important light on the amount of color perceived in a snapshot of experience.

Another important question raised by these results is related to the relationship between our findings and the well-established limits of peripheral vision. Specifically, can our findings simply be accounted for by factors such as visual crowding or the drop off in retinal acuity (Pelli & Tillman, 2008; Whitney & Levi, 2011; Rosenholtz, 2016)? We believe that while those bottlenecks on perception are thematically related to the results reported here, they do not constrain perception nearly enough to explain our results. Perhaps the clearest bit of data to explain this idea comes from Wallis et al. (2019) who directly explored the boundary limits of perception in the periphery using the same texture synthesis algorithm as we did in Experiments 1 and 3 (Freeman & Simoncelli, 2011). In their study, Wallis and colleagues explicitly asked participants to determine which image in a set had been synthesized (i.e., an odd-ball detection task). In this case, the critical image had been scrambled in a manner that corresponds to our small scramble condition in Experiment 1. In that case, observers were able to reliably identify the synthesized image amongst natural images with little difficulty. As a reminder, in our Experiment 1, not a single person noticed the critical stimuli in the small scramble condition. Thus, it seems then as if our results do not stem from a fundamental limit on peripheral vision imposed by things like crowding or the drop off in retinal acuity; instead, we argue that the primary bottleneck is likely the way in which observers are deploying their attention across these scenes. An interesting line of future research would be to explore exactly how attention is generally deployed in natural viewing conditions and what the exact resolution of perceptual awareness is in those situations.

Finally, the idea of the resolution of perception is intimately related to a similar idea in the literature referred to as the useful field of view (UFOV; Mackworth, 1965). The UFOV is the amount of visual information that can be processed within a single fixation without any eye or head movements. Critically, however, is the fact that UFOV appears to vary as a

function of attentional load and task complexity (Chan & Courtney, 1998; Atchley & Dressel, 2004; Ringer et al., 2016), much in the same way that we suggest the resolution of perception varies as a function of attentional allocation. In other words, as the attentional load increases or gets more focused around the fovea, the UFOV shrinks and the amount of information observers can usefully process decreases. One interesting idea for future research raised by research on the UFOV is the notion that as people age, the resolution of perception varies. It has been systematically shown that the UFOV varies as a function of age (Ball et al., 1988; 1993; Sekuler et al., 2000). These findings raise the interesting possibility that older and younger adults may have different rates of inattention blindness and may actually differ in terms of how much they can perceive “in the blink of an eye.”

Overall, our results reveal that observers can gaze upon an image, process enough of that image to identify it out of multiple options, but completely fail to notice drastic changes to the periphery of that image. These results raise a variety of experimentally tractable questions to further understand the richness of experience in terms of both an individual snapshot as well as when naturally embedded in the real world.

Acknowledgements

This study was supported by a National Science Foundation Collaborative Research Award (#1829470) to M.A.C. Thanks to Rachel Albert, Thomas Botch, Sarah Cormiea, Leyla Isik, Nancy Kanwisher, Ken Nakayama, and Caroline Robertson for extensive discussion and feedback on the project. Thanks to Sebastian Montesinos for help with data collection.

References

- Atchley, P., & Dressel, J. (2004). Conversation limits the functional field of view. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 46(4), 664–673.
- Baars, B.J. (2002) The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6, 47-52.
- Ball, K. K., Beard, B. L., Roenker, D. L., Miller, R. L., & Griggs, D. S. (1988). Age and visual search: Expanding the useful field of view. *Journal of the Optical Society of America*, 5, 2210–2219.
- Ball, K. K., Owsley, C., Sloane, M. E., Roenker, D. L., & Bruni, J. R. (1993). Visual attention problems as a predictor of vehicle crashes in older drivers. *Investigative Ophthalmology & Visual Science*, 34, 3110–3123.
- Block, N. (2007) Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30, 481-5478.
- Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 16, 567-575.

Boly, M., Massimini, M., Tsuchiya, N., Postle, B.R., Koch, C., & Tononi, G. (2017). Are the neural correlates of consciousness in the front or in the back of the cerebral cortex? Clinical and neuroimaging evidence. *Journal of Neuroscience*, 37, 9603-9613.

Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision* 10, 433-436.

Castelhano, M.S., & Henderson, J.M. (2008) The influence of color on scene gist. *Journal of Experimental Psychology: Human Perception and Performance*, 34, 660-675.

Castelhano, M.S., & Heaven, C. (2011). Scene context influences without scene gist: Eye movements guided by spatial associations in visual search. *Psychonomic Bulletin and Review*, 18, 890-896.

Chan, H. S., & Courtney, A. J. (1998). Stimulus size scaling and foveal load as determinants of peripheral target detection. *Ergonomics*, 41, 1433–1452.

Cohen, M.A., Alvarez, G.A., & Nakayama, K. (2011). Natural- scene perception requires attention. *Psychological Science*. 22 (9) 1165-1172

Cohen, M.A., Cavanagh, P., Chun, M.M., & Nakayama, K. (2012). The attentional requirements of consciousness. *Trends in Cognitive Sciences*, 16, 411-417.

Cohen, M.A., Dennett, D.C., & Kaniwsher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences*, 20, 324-335.

Dagnelie, G. (2011). *Visual Prosthetics: Physiology, Bioengineering, & Rehabilitation*. New York, NY: Springer.

Dehaene, S., & Changeux, J.P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*. 70, 220-227.

Dennett, D.C. (1991) *Consciousness Explained*, Little Brown: Cambridge, USA.

Deubel, H. & Schneider, W.X. (1996) Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Research*.

Fei-Fei, L., Iyer, A., Koch, C., & Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of Vision*, 7(1), 10:1–29

Freeman, J., & Simoncelli, E.P. (2011). Metamers of the ventral stream. *Nature Neuroscience*, 14, 1195-1201.

Frisén, L. (1990). Clinical tests of vision. New York, NY: Raven Books.

Greene, M. R., & Oliva, A. (2009). The briefest of glances: The time course of natural scene understanding. *Psychological Science*, 20, 464–472.

Haun, A.M., Tononi, G., Koch, C., Tsuchiya, N. (2017). Are we underestimating the richness of visual experience? *Neuroscience of Consciousness*, doi: 10.1093/nc/niw023.

Hayhoe, M.M., Shrivastava, A., Mruczek, R., & Pelz, J. (2003). Visual memory and motor planning in a natural task. *Journal of Vision*, 3, 49-63.

Henderson, J.M. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, 7, 498-504.

Hoffman, J.E. & Subramaniam, B. (1995) The role of visual attention in saccadic eye movements. *Perception and Psychophysics*, 57, 787-795.

Hollingworth, A. (2009). Memory for real-world scenes. In J. R. Brockmole (Ed.), *The Visual World in Memory* (pp. 89-119). Hove, UK: Psychology Press.

Intraub, H. (1981). Rapid conceptual identification of sequentially presented pictures. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 604–610.

Knotts, J.D., Odegaard, B., Lau, H., & Rosenthal, D. (2019). Subjective inflation: phenomenology's get-rich-quick scheme. *Current Opinion in Psychology*, 29, 49-55.

Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, 11, 16-22.

Komatsu, H. (2006). The neural mechanisms of perceptual filling-in. *Nature Reviews Neuroscience*, 7, 220-231.

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14, 301-307.

Kowler, E., Anderson, E., Doshier, B., & Blaser, E. (1995) The role of attention in the programming of saccades. *Vision Research*, 35, 1897-1916.

Lamme, V. A. F. (2003) Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7, 12-18.

Lamme, V.A.F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, 1, 204-220.

Loschky, L. C. & McConkie, G. W. (2002). Investigating spatial vision and dynamic attentional selection using a gaze-contingent multiresolutional display. *Journal of Experimental Psychology: Applied*, 8, 99–117.

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

Mackworth, N. H. (1965). Visual noise causes tunnel vision. *Psychonomic Science*, 3, 67–68.

McConkie, G.W. & Rayner, K. (1976) Asymmetry of the perceptual span in reading. *Bulletin of the Psychonomic Society*. 8, 365-368.

- Mielliet, S., Zhou, X., He, L., Rodger, H., & Caldara, R. (2010) Investigating cultural diversity for extrafoveal information use in visual scenes. *Journal of Vision*, 10, 1-18.
- Millodot, M. (2014). *Dictionary of Optometry and Visual Science*. Elsevier Health Sciences: UK.
- Nuthmann, A. (2017). Fixation durations in scene viewing: Modeling the effects of local image features, oculomotor parameters, and task. *Psychonomic Bulletin and Review*, 24, 370-392.
- Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Pelli, D. G., & Tillman, K. A. (2008) The uncrowded window of object recognition. *Nature Neuroscience*, 11, 1129-1135
- Pelz, J.G., & Canosa, R. (2001). Oculomotor behavior and perceptual strategies in complex tasks. *Vision Research*, 41, 3587-3596.
- Polyak, S.L. (1941). *The Retina*. University of Chicago Press: Chicago.
- Pomplun, M., Reingold, E.M., Shen, J. (2001) Peripheral and parafoveal cueing and masking effects on saccadic selectivity in a gaze-contingent window paradigm. *Vision Research*, 41, 2757-2769.
- Portilla, J., & Simoncelli, E.P. (2000). A parametric texture model based on joint statistics of complex wavelet coefficients. *International Journal of Computer Vision*, 40, 49-71.
- Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3-25.
- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 509–522.
- Rahnev, D., Maniscalco, B., Graves, T., Huang, E., de Lange, F.P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, 14, 1513-1515.
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological Bulletin*, 124, 372-422.
- Rayner, K., Smith, T.J., Malcolm, G.L., & Henderson, J.M. (2008). Eye movements and visual encoding during scene perception. *Psychological Science*, 20, 6-10.
- Rayner, K. (2014) The gaze-contingent moving window in reading: Development and review. *Visual Cognition*, 22, 242-258.

- Reingold, E., M., Charness, N., Pomplun, M., & Stampe, D. M. (2001). Visual span in expert chess players: evidence from eye movements. *Psychological Science*, 12, 48–55.
- Rensink, R.A. (2000) The dynamic representation of scenes. *Visual Cognition*, 7, 17-42.
- Rensink, R., 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*, 8, 368-373.
- Ringer, R. V., Throneburg, Z., Johnson, A. P., Kramer, A. F., & Loschky, L. C. (2016). Impairing the useful field of view in natural scenes: Tunnel vision versus general interference. *Journal of Vision*, 16,1-25.
- Rönne, H. (1915). Zur Theorie und Technik der Bjerrumschen Gesichtsfelduntersuchung. *Archiv für Augenheilkunde*, 78(4), 284–301.
- Rosenholtz, R. Capabilities and limitations of peripheral vision. (2016). *Annual Review of Vision Science*, 2, 437-457.
- Sampanes, A.C., Tseng, P., & Bridgeman, B. (2008) The role of gist in scene recognition. *Vision Research*, 48, 2275-2283.
- Savino, P.J., Danesh-Meyer, H. V. (2012). Color Atlas and Synopsis of Clinical Ophthalmology -- Wills Eye Institute -- Neuro-Ophthalmology. Lippincott Williams & Wilkins
- Sekuler, A. B., Bennett, P. J., & Mamelak, M. (2000). Effects of aging on the useful field of view. *Experimental Aging Research*, 26, 103–120.
- Simons, D., & Levin, D. (1997). Change blindness. *Trends in Cognitive Sciences*, 1, 261–267.
- Simons, D.J., & Rensink, R. (2005). Change blindness: Past, present, and future. *Trends in Cognitive Sciences*, 9, 16-20.
- Solovey, G., Graney, G.G., & Lau, H. (2015). A decisional account of subjective inflation of visual perception at the periphery. *Attention Perception and Psychophysics*, 77, 258-271.
- Swanson, W.H., & Fish, G.E. (1995). Color matches in diseased eyes with good acuity: Detection of deficits in cone optical density and in chromatic discrimination. *Journal of the Optical Society of America A*. 12, 2230-2236.
- Tatler, B., Gilchrist, I., & Risted, J. (2003). The time course of abstract visual representation. *Perception*, 32, 579–593
- Tononi G, Boly M, Massimini M, Koch C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience* 17, 450– 461.
- Torralbo, A., Walther, D. B., Chai, B., Caddigan, E., Fei-Fei, L., & Beck, D. M. (2013). Good exemplars of natural scene categories elicit clearer patterns than bad exemplars but not greater BOLD activity. *PLoS ONE*, 8,e58594

Tyler, C.W. (2014). Peripheral color demo. *iPerception*. 6, 1-5.

Unema, P.J.A., Pannasch, S., Joos, M., & Velichkovsky, B.M. (2005). Time course of information processing during scene perception: The relationship between saccade amplitude and fixation duration. *Visual Cognition*, 12, 473-494.

Van Boxtel, J.A., Tsuchiya, N., & Koch, C. (2010) Consciousness and attention: on sufficiency and necessity. *Frontiers in Psychology*, 1, 1-13

Võ, M.L.-H. & Henderson, J.M. (2011) Object-scene inconsistencies do not capture gaze: evidence from the flash-preview moving-window paradigm. *Attention Perception, and Psychophysics*, 73, 1742-1753.

Wallis, T.S.A., Funke, C.M., Ekcer, A.S., Gatys, L.A., Wichmann, F.A., & Bethge, M. (2019). Image content is more important than Bouma's Law for scene metamers. *eLife*. e42512.

Walther, D. B., Caddigan, E., Fei-Fei, L., & Beck, D. M. (2009). Natural scene categories revealed in distributed patterns of activity in the human brain. *Journal of Neuroscience*, 29, 10573–10581.

Whitney, D. & Levi, D.M. (2011) Visual crowding: a fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences*, 15, 160-168.

Appendix

Reference	Fixation duration measurement
Rayner, 1998	275 ms
Pelz & Canosa, 2001	327 ms
Hayhoe et al., 2003	225 ms
Henderson, 2003	330 ms
Unema et al., 2005	260 ms
Rayner et al., 2008	256 ms
Castelhano & Heaven, 1998	260 ms
Nuthmann, 2017	275 ms

Table A1. References to prior studies that measured the duration of fixations when gazing upon a variety of stimuli and performing a variety of tasks. The average of these measurements were used to determine how long stimuli were presented for across all experiments.

¹ Since the full visual field extends roughly 110° to the left and right of fixation (i.e., temporally) and roughly 60° above fixation and 70°-75° below fixation, our preserved circular aperture made up an incredibly small portion of the visual field (Rönne, 1915; Frisén, 1990; Dagnelie, 2011; Savino & Danehs-Meyer, 2012).