Familiarity alters the bandwidth of perceptual awareness

Michael A. Cohen[1,2], Skyler Sung[2], & Zaki Alaoui[2]

[1] McGovern Institute for Brain Research, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology
[2] Department of Psychology and Program in Neuroscience, Amherst College

Abstract

Results from paradigms such as change blindness and inattentional blindness indicate that observers are unaware of numerous aspects of the visual world. However, intuition suggests that perceptual experience is richer than these results indicate. Why does it feel like we see so much when the data suggests we see so little? One possibility stems from the fact that experimental studies typically present observers with stimuli that they have never seen before. Meanwhile, when forming intuitions about perceptual experience, observers reflect on their experiences with scenes with which they are highly familiar (e.g., their office).Does prior experience with a scene change the bandwidth of perceptual awareness? Here, we asked if observers were better at noticing alterations to the periphery in familiar scenes compared to unfamiliar scenes. Here, we found that observers noticed changes to the periphery more frequently with familiar stimuli. Signal detection theoretic analyses revealed that when observers are unfamiliar with a stimulus, they are less sensitive at noticing peripheral alterations (*d’*) and are more conservative in their response criterion (*c*). Taken together, these results suggest that prior knowledge expands the bandwidth of perceptual awareness. It should be stressed, that these results challenge the widely held idea that prior knowledge fills-in perception. Overall, these findings highlight how prior knowledge plays an important role in determining the limits of perceptual experience and is an important factor to consider when attempting to reconcile the tension between empirical observation and personal introspection.

Introduction

How much of the visual world are observers aware of at any given moment? Over the last few decades, several results from paradigms like change blindness and inattentional blindness suggest that observers fail to consciously perceive many items in the world around them (Rensink et al., 1997; Mack & Rock, 1997; O'Regan et al., 1999; Simons et al., 1999; Rensink, 2002; Most et al., 2005; Jensen et al., 2011). For example, researchers have found that observers routinely fail to notice a fight occurring right in front of them (Chabris et al., 2011), a gorilla embedded in the computed tomography (CT) scan of a patient's lungs (Drew et al., 2013), and salient items on a road that can result in a car accident (Most & Astur, 2007). Taken together, these results have been cited to argue that observers are aware of surprisingly little of the visual world and perception is quite sparse (Kouider et al., 2010; Cohen et al., 2011; 2012; Dehaene, 2014; Odegaard et al., 2018; Knotts, et al., 2019).

Although results such as these have been thoroughly studied and replicated numerous times, there is one longstanding lingering issue related to this work: virtually no one believes their perceptual experience is as impoverished as these results suggest (Koch & Tsuchiya, 2007; Block, 2011; Haun et al., 2017; Lamme, 2018). Indeed, part of what makes change blindness and inattentional blindness so compelling is that observers intuitively believe that they would easily be able to spot these salient changes and events (Levin et al., 2000; Scholl et al., 2004; Beck et al., 2007; Loussouarn et al., 2011). Thus, the question remains, why do observers believe they see so much when the data suggests they see so little?

One potential factor leading to the discrepancy between empirical observation and personal introspection is based on the stimuli used in these experiments. In every study examining the limits of perceptual awareness with natural scenes, observers are presented with stimuli they have never seen before. While observers may have no problem understanding, say, a picture of an airplane on a runway, they will have never previously seen *that* particular picture of *that* specific airplane on *that* particular runway. Meanwhile, when introspecting on the richness of their perceptual experiences, observers will consider stimuli with which they are highly familiar or currently occupy (e.g., their own office or kitchen). Therefore, there is a systematic discrepancy between the stimuli used to test the limits of perceptual awareness (i.e., unfamiliar) and the types of situations observers consider when thinking about the richness of their experience (i.e., familiar).

Does familiarity with a scene change the bandwidth of perceptual experience? Perhaps the most common answer to this question is that prior knowledge about a scene will give observers a richer perceptual experience since that knowledge allows the brain to fill-in specific details of the visual world even though those details are not physically present (Komatsu, 2006). The most well-studied example of filling-in is in relation to the physiological blind spot, a region of the retina that has no photoreceptors and yet observers are unaware of this missing input to the visual system because the brain fills in the missing detail using the surrounding context (Ramachandran & Gregory, 1991). Prior research has shown that many different elements, such as brightness, color, texture, motion, and spatiotemporal information can nonetheless be perceived at regions of space corresponding to the blind-spot (Ramachandran,1992; Tripathy et al., 1996; Pessoa et al., 1998; Maus & Whitney, 2016). Although filling-in has been primarily shown with lower-level features and is associated with early visual cortex (e.g., V1, V2, etc. Matsumoto & Komatsu, 2005; Komatsu, 2006; De Weerd et al., 2006; Weil & Rees, 2011), it is widely claimed that a similar process may extend to higher-lever perception and provide observers with the impression of a rich, detailed visual experience (Lau & Rosenthal, 2011; Chong et al., 2016; Otten et al., 2017; Tyler & Solomon, 2018; Toscani et al., 2021; Seth, 2021). To provide a concrete example of how this process might work, consider previous studies showing that observers routinely fail to notice when color is entirely removed from the periphery (Balas & Sinha, 2007; Cohen & Rubenstein, 2020; Cohen et al., 2020). Why do observers feel as if they see color across their visual field even though these results suggest that they do not? According to the filling-in hypothesis, observers' familiarity with their surroundings allows items in the periphery to appear colorful because prior knowledge fills-in those colors, providing the impression of a rich perceptual experience.

An alternative possibility is that prior knowledge about a scene will expand the limits of perceptual awareness and allow observers to notice certain details more easily across the visual field. Consider again those studies suggesting that observers are often unaware of color in the periphery. Under this view, the frequency with which observers notice that color has been removed from the periphery would vary as a function of prior exposure to the scene. When observers are unfamiliar with a scene, noticing that, say, a red stop sign and a yellow school bus are desaturated in the periphery would be relatively difficult. However, when observers are familiar with the scene, noticing that those items are desaturated in the periphery would be easier. The key idea here is that when observers have a pre-established expectation about how the world should look, they will be better able to notice when the sensory input does not meet those expectations. Indeed, prior work using inattentional blindness paradigms have shown how participants' expectations affect what they consciously perceive (Most et al., 2004; Ward & Scholl, 2014). Thus, under this framework, participants' ability to notice changes to their visual environment will expand.

To differentiate between these two possibilities, we measured the outer limits of peripheral vision by presenting stimuli across the entire horizontal plane in a custom-built display dome (Figure 1). Specifically, in Experiment 1, we used a staircase procedure to identify the threshold at which observers were at chance (i.e., ~50%) at determining whether color had been removed from the periphery or the periphery had been scrambled. In Experiment 2, we used signal detection theoretic analyses to examine how the perceptual sensitivity (*d'*) or response bias (*c*) change as a function of familiarity.

In this case, the filling-in and expansion hypotheses make two opposite predictions. Consider again a case in which color is removed from the periphery. According to the filling-in hypothesis, when observers are shown familiar stimuli, they will consistently report that the periphery is in color when it is actually desaturated. Therefore, under this view, realizing the periphery is actually desaturated will be *harder* in the familiar condition relative to the unfamiliar condition. For this reason, to find the ~50% threshold, the size of the unaltered region will have to become *smaller*. Conversely, according to the expansion hypothesis, prior experience with the stimulus allows observers to form a mental representation of the original, unaltered scene that they can then compare to the test image. Therefore, noticing that the periphery is desaturated will be *easier* in the familiar condition relative to the unfamiliar condition. For this reason, to find the ~50% threshold, the size of the unaltered region will have to become *larger.*

To preview the results, we found support for the expansion hypothesis: observers were significantly better at detecting when the periphery was desaturated (Experiment 1a) or scrambled (Experiment 1b) when they were familiar with the stimulus compared to when they were not familiar with the stimulus. Moreover, signal detection theoretic analyses indicated that this effect was driven by both a change in perceptual sensitivity and response bias. When observers were unfamiliar with a stimulus, they were less sensitive to changes in the periphery and employed a more conservative decision criterion. These results help reconcile the tension between empirical findings showing that perception is surprisingly sparse and the intuition that perception is quite rich; our intuitions are developed in familiar settings while our perceptual abilities are measured with unfamiliar stimuli and this discrepancy in familiarity leads to a tension between empirical measures and subjective intuitions. Together, these results suggest that paradigms like change blindness and inattentional blindness may systematically underestimate the richness of experience.

Experiment 1: Peripheral threshold

Methods

Preregistration: All aspects of the experiment — including stimulus creation, experimental procedures, number of participants, and planned analyses — were preregistered on the Open Science Framework: https://osf.io/9y6nw/. In addition, all stimuli and raw data are available on the Open Science Framework.

Participants: For both Experiments 1a and 1b, we preregistered a sample size of 20 participants for both experiments. This sample size was determined in advance based on pilot testing and then pre-registered before official data collection began. All participants were recruited from the Amherst College community and had normal or corrected-to-normal vision, no known neurological conditions, and passed the Ishihara color test (Birch, 1997). The experimental procedures were approved by the Institutional Review Board at Amherst College and informed consent was obtained from each participant, who were compensated for their time.

Stimuli: The experiments were performed on an ImmersaVu Desktop Dome (Figure 1). This display uses a projection system that reflects onto a glass fiber composite, which was assembled to create a 180° field of view across the horizontal axis (and 67.5° of visual angle across the vertical axis). Participants were seated approximately 76 cm from the monitor during the experiments and their head was stabilized on a chin rest. All stimuli were presented using MATLAB (2022a) with the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997).

Full-field viewing display

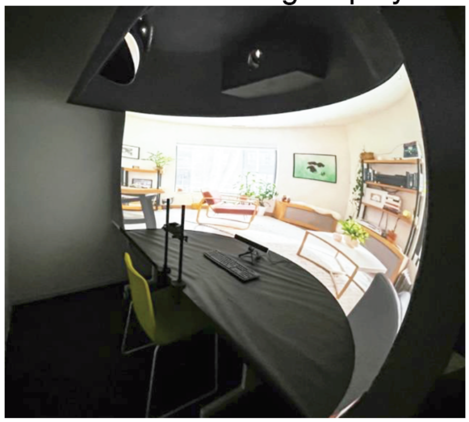

Figure 1. Picture of the full field viewing display used in both experiments. Note the chin rest, the eye-tracker, and how the display wraps around to the ends of the table (i.e., 180° of visual angle).

The stimulus set comprised 186 images that were obtained on the internet and via the experimenter's private photograph collection. Each image was 3840x2160 pixels, which is the resolution of the ImmersaVu display, and spanned the entire display. The stimulus set was made up of a wide variety of images of both indoor and outdoor natural scenes: gymnasiums, concert halls, atriums, grocery stores, restaurants, forests, beaches, mountain ranges, lakes, street scenes, highways, and so forth (see OSF link for full stimulus set).

In Experiment 1a, the center part of the image was in color, while the periphery was completely desaturated. The color portion was preserved by a circular flat top Gaussian that varied in size as a function of participants' performance equally along the horizontal and vertical axes (Figure 2A). For Experiment 1b, the center part of the image was completely desaturated and unscrambled, while the periphery was scrambled (Figure 2B). To scramble the images, we modified a texture synthesis algorithm developed by Freeman & Simoncelli (2011) by tiling each individual image with 200 evenly sized/spaced pooling windows in a 20x10 grid. As was the case in Experiment 1a, the center portion was preserved by a flat top Gaussian that varied in size as a function of participants' performance.

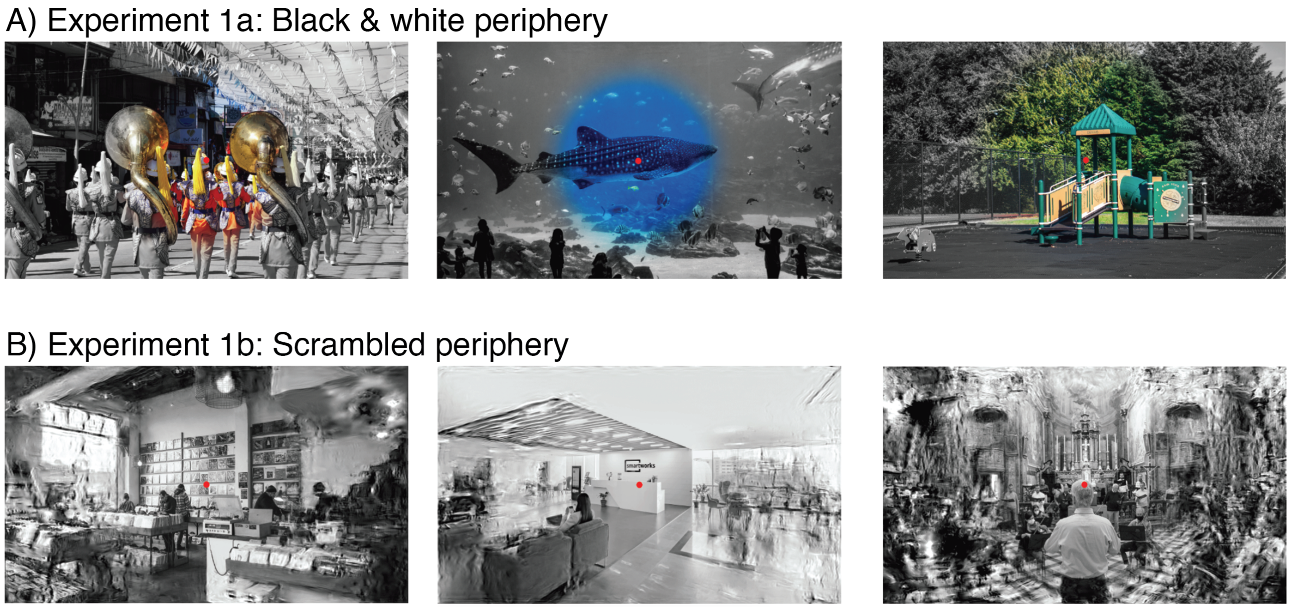

Figure 2. Sample stimuli and results from Experiments 1a (A) and 1b (B). For each experiment, example stimuli are shown during the thresholding process (though note that when presented in the experiment, they were presented on curved displays that took up of the participants' entire visual field along the horizontal axis).

Procedures: The task participants performed was determining whether an image had been altered in some way in the periphery or not. In Experiment 1a, the task was to report if the entire image was a) entirely in color or b) if only the center was in color and the periphery was desaturated (i.e., "black and white in the periphery"). In Experiment 1b, the task was to report if a) the entire image was entirely unscrambled (i.e., normal) or b) if only the center was unscrambled and the periphery was scrambled. Thus, on each trial, participants reported what they believed the image looked like on that trial.

There were two main experimental conditions: 1) unfamiliar and 2) familiar (Figure 3). On the unfamiliar trials, the participant pressed the space bar to start the trial. Once they pressed the space bar, a blank display appeared with nothing but a fixation dot for 500ms. After 500ms, the test image was presented on the display for 283ms while participants maintained fixation on a central dot. Critically, the participant had no prior knowledge of what the image would be in advance of it being shown. On the familiar trials, however, participants were shown a full color version of the test image (Experiment 1a) or a fully unscrambled desaturated image (Experiment 1b) before seeing the test image. When that study image was shown, it remained on the screen for a minimum of 6 seconds while participants inspected it in whatever way they chose. After at least 6 seconds, the trial began whenever participants wanted by pressing the space bar. Once they pressed the space bar, the study image disappeared, and a blank display appeared with nothing but a fixation dot for 500ms. After 500ms, the test image was shown for 283ms.

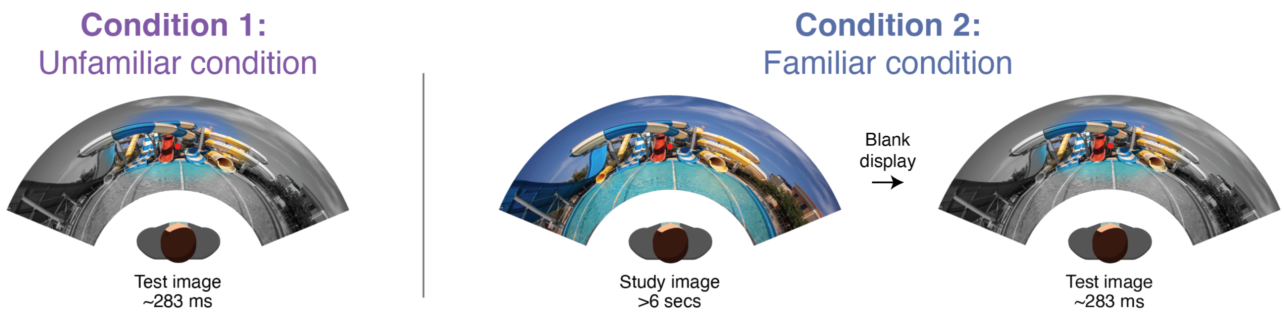

Figure 3. Visualization of one trial for the two experimental conditions. Left panel: Unfamiliar condition in which participants are shown the test image for approximately 283ms. Right panel: familiar condition in which participants are shown a study image that has not been altered for a minimum of 6 seconds before being shown the altered test image for approximately 283ms.

The experiment was broken up into four blocks: 2 blocks for the unfamiliar condition and 2 blocks for the familiar condition. Each block was comprised of 60 trials (240 total trials across all four blocks). For half the participants, the order in which the blocks were presented was familiar, unfamiliar, familiar, and unfamiliar. For the other half of participants, the order was unfamiliar, familiar, unfamiliar, and familiar. For each block, we determined the ~50% threshold in terms of degree of visual angle at which participants could not differentiate whether the test image had been altered (i.e., color center and desaturated periphery in Experiment 1a, or unscrambled center and scrambled periphery in Experiment 1b). At the end of the experiment, four threshold values were obtained (2 for unfamiliar and 2 for familiar). These two threshold values were averaged together within each participant to obtain one final threshold value.

To determine this threshold, we used QUEST, a Bayesian adaptive staircase procedure (Watson and Pelli, 1983). Although participants were given two options, in reality, the test image was always altered in the periphery in terms of either being desaturated or scrambled. However, by virtue of QUEST changing the size of the center region (i.e., color in Experiment 1a and unscrambled in Experiment 1b) participants were unaware of this as the algorithm found the point at which they said altered/unaltered ~50% of the time. Participants were unaware of the fact that every image was altered and were told that some images would be unaltered across the visual field and others would be altered in some form or another.

The experiment began with an eye-tracking calibration sequence. The eye-tracker was used to ensure fixation was maintained throughout every trial. Here, we used a Gazepoint GP3 eye tracker [https://www.gazept.com/product/gazepoint-gp3-eye-tracker/]. Before each block, participants performed a nine-point calibration sequence that required 1° of visual angle to be accepted. Once each trial began, if participants broke fixation from a fixation dot in the middle of the image (i.e., more than 2° of visual angle away from fixation), the trial was discarded and not included in the QUEST procedure. At the beginning of the first and third blocks, participants performed a recalibration sequence to ensure quality eye-tracking and minimize the influence of drift during the experiment.

Results

Possible experimental outcomes: Broadly speaking, there are two main hypotheses for these experiments.

Hypothesis 1 – The filling-in hypothesis: According to this hypothesis, prior exposure to the target image will allow the visual system fill-in to certain details in the periphery that are not physically present on the display (Lau & Rosenthal, 2011; Chong et al., 2016; Otten et al., 2017; Tyler & Solomon, 2018; Toscani et al., 2021; Seth, 2021). Therefore, accurately determining whether the periphery is, say, colorful or desaturated will be *harder*. This is because observers will say color is present, even when it is not, because they are seeing colors that have been filled-in by their visual system. Thus, in order to find the ~50% threshold, the colored part of the display will have to be made *smaller* to make the task somewhat easier (Figure 4, top right).

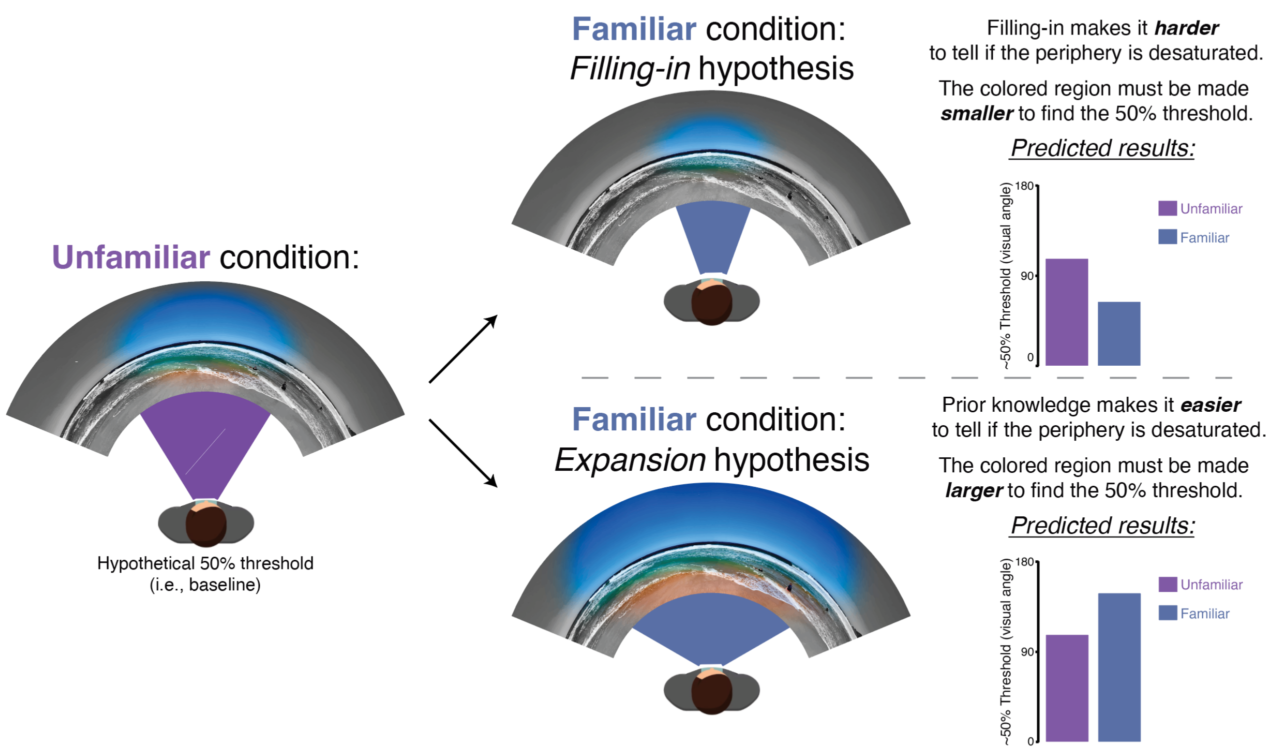

Figure 4. Potential experimental outcomes with the current design. Left panel represents the hypothetical ~50% threshold when participants are unfamiliar with the stimuli. This can be thought of as the baseline condition. Top right panel represents prediction made by the filling-in hypothesis. Bottom right panel represents the prediction made by the expansion hypothesis.

Hypothesis 2 – The expansion hypothesis: According to this hypothesis, prior knowledge will allow participants to form a mental representation of what the unaltered, study image looks like. Then, they can directly compare that internal representation to the sensory signal, making it *easier* to determine if, say, the periphery is colorful or desaturated (e.g., "I know the man's shirt should be bright red, but they are not, so the periphery must be desaturated."). Therefore, to find the ~50% threshold, the colored part of the display will have to be made *larger* to make the task somewhat harder (Figure 4, bottom right).

Experiment 1a results (Color): Overall, the results were unambiguous. We found a significant difference in the ~50% threshold for detecting color vs. no color when comparing the unfamiliar condition with the familiar condition (Figure 5A). Specifically, the outer threshold was significantly larger in the familiar condition (142° of visual angle) than in the unfamiliar condition (111° of visual angle; $t(19)=4.99$, $p=0.00009$). This result lends support to the hypothesis that prior knowledge expands perception and challenges the hypothesis that such knowledge may potentially fill-in perception.

Experiment 1b results (Scrambled): Does the same result found with color also extend to other alterations to the periphery? When scrambling the periphery, we once again found a significant difference in the ~50% threshold for detecting scrambled vs. unscrambled when comparing the unfamiliar condition with the familiar condition (Figure 5B). Specifically, the outer threshold was significantly larger in in the familiar condition (122° of visual angle) than in the unfamiliar condition (99° of visual angle; $t(19)=5.89$, $p=0.00001$).

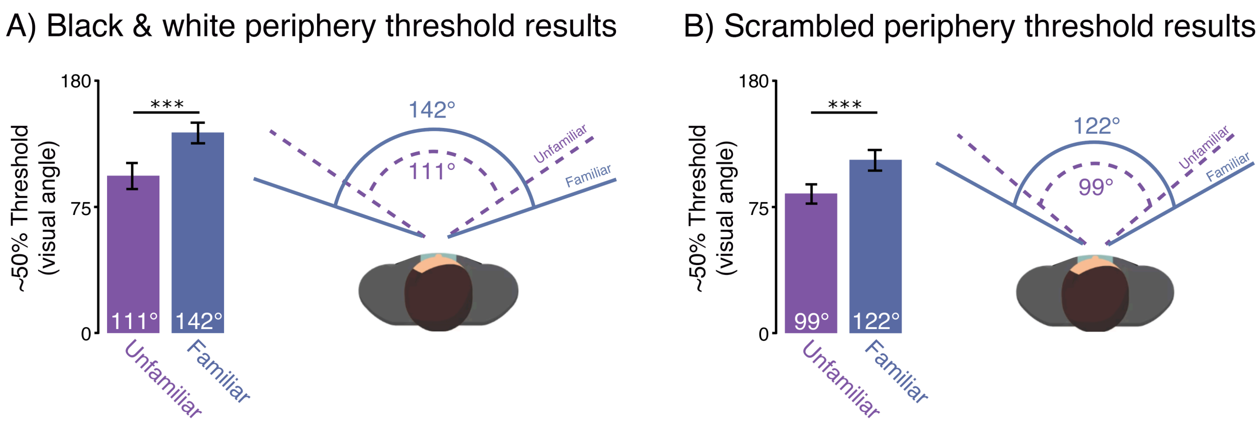

Figure 5. Experiment 1 results. A) Experiment 1a (black and white periphery). B) Experiment 1b (scrambled periphery. The y-axis represents the output of the QUEST procedure in identifying the ~50% threshold at which participants were at chance in determining if the image has been altered in the periphery or not. Purple bars represent the unfamiliar condition and blue bars represent the familiar condition. Error bars represent the standard error of the mean. In addition, diagrams visualize the size of the unfamiliar and familiar thresholds relative to an observer's field of view. ****p*<0.001.

Discussion

Overall, we found that with several types of alterations to the periphery, participants were better at detecting those alterations when a stimulus was familiar compared to when it was unfamiliar. These results are consistent with predictions made by the expansion hypothesis and are inconsistent with predictions made by the filling-in hypothesis.

Experiment 2: Signal detection analysis

What is driving the effects in Experiment 1? Why are observers better at detecting peripheral alterations with familiar stimuli? Although we cannot provide full answers to these questions (see Discussion), the goal of Experiment 2 was to make some progress in this domain. Specifically, Experiment 2 was designed such that we could perform signal detection theoretic analyses to determine if either perceptual sensitivity (*d'*) or response bias (*c*) change as a function of familiarity (MacMillan & Creelman, 1991; Rahnev et al., 2011). In this case, a change in sensitivity (*d'*) would indicate a change in participants' ability to discriminate between signal and noise when attempting to detect peripheral alterations. Meanwhile, a change in criterion (*c*) would indicate that participants are overall more likely to say the periphery has (or has not) been altered.

Methods

Preregistration: All aspects of the experiment — including stimulus creation, experimental procedures, number of participants, and planned analyses — were preregistered on the Open Science Framework: https://osf.io/9y6nw/. In addition, all stimuli and raw data are available on the Open Science Framework.

Participants: For both Experiments 2a and 2b, we preregistered a sample size of 10 participants who would perform both experimental conditions. This sample size was determined in advance based on pilot testing and then pre-registered before official data collection began. All participants were recruited from the Amherst College community and had normal or corrected-to-normal vision, no known neurological conditions, and passed the Ishihara color test (Birch, 1997). The experimental procedures were approved by the Institutional Review Board at Amherst College and informed consent was obtained from each participant, who were compensated for their time or received course credit.

Stimuli: The same stimuli used in Experiment 1 were used in Experiment 2. In this case, however, the size of the central preserved region did not change on the trials in which the periphery was altered. In the color condition (Experiment 2a), the size of the preserved region was 90° of visual angle, while in the scrambled condition (Experiment 2b), the size of the preserved region was 85° of visual angle.

Procedures: Participants performed 4 blocks of 100 trials (1) color familiar, 2) color unfamiliar, 3) scrambled familiar, and 4) scrambled unfamiliar). The order in which each block was completed was counterbalanced across participants with the only constraint that the two stimuli conditions were performed back-to-back (i.e., color familiar followed by color unfamiliar, etc.). Within each block, half the trials were altered in the periphery and half were not. The order in which altered/unaltered trials were presented was randomized within each block. On the unfamiliar trials, the participant pressed the space bar to start the trial. Once they pressed the space bar, a blank display appeared with nothing but a fixation dot for 500ms. After 500ms, the test image was presented on the display for 283ms while participants maintained fixation on a central dot. Critically, the participant had no prior knowledge of what the image would be in advance of it being shown. On the familiar trials, however, participants were shown a full color version of the test image (Experiment 2a) or a fully unscrambled desaturated image (Experiment 2b) before seeing the test image. When that study image was shown, it remained on the screen for a minimum of 6 seconds while participants inspected it in whatever way they chose. After 6 seconds, the trial began when participants pressed the space bar. Once they pressed the space bar, the study image disappeared, and a blank display appeared with nothing on it but a fixation dot for 500ms. After 500ms, the test image was shown for 283ms.

As was the case in Experiment 1, the experiment began with an eye-tracking calibration sequence. The eye-tracker was used to ensure fixation was maintained throughout every trial. Once again, we used a Gazepoint GP3 eye tracker. Before each block, participants performed a nine-point calibration sequence that required 1° of visual angle to be accepted. Once each trial began, if participants broke fixation from a fixation dot in the middle of the image (i.e., more than 2° of visual angle away from fixation), the trial was discarded and not included in the QUEST procedure. At the beginning of the first and third blocks, participants performed a recalibration sequence to ensure quality eye-tracking and minimize the influence of drift during the experiment.

Results

Experiment 2a results (color): Overall, we found that observers were more accurate in the familiar condition relative to the unfamiliar condition (familiar accuracy = 86%; unfamiliar accuracy = 71%; $t(9)=4.23$, $p=0.002$). In addition, d' was higher in the familiar condition relative to the unfamiliar condition (familiar $d'=2.32$; unfamiliar $d'=1.28$; $t(9)=3.76$, $p=0.004$), while c was lower in the familiar condition relative to the unfamiliar conditions (familiar $c=0.08$; unfamiliar $c=0.34$; $t(9)=2.61$, $p=0.03$). It is also worth noting that in the familiar condition, c was not significantly greater than zero ($t(9)=0.84$, $p=0.42$), indicating that observers were equally likely to say a trial had or had not been altered. Conversely, in the unfamiliar condition, c was significantly greater than zero ($t(9)=3.84$, $p=0.004$), indicating that when observers are not familiar with a stimulus, they are predisposed to say that no alteration was made.

Experiment 2b results (scrambled): Overall, we found that observers were more accurate in the familiar condition relative to the unfamiliar condition (familiar accuracy = 88%; unfamiliar accuracy = 73%; $t(9)=4.40$, $p=0.002$). In addition, d' was higher in the familiar condition relative to the unfamiliar condition (familiar $d'=2.68$; unfamiliar $d'=1.53$; $t(9)=3.49$, $p=0.007$), while c was lower in the familiar condition relative to the unfamiliar conditions (familiar $c=0.19$; unfamiliar $c=0.55$; $t(9)=2.84$, $p=0.02$). It is also worth noting that in the familiar condition, c was trending, but not significantly greater than zero ($t(9)=1.9$, $p=0.09$), indicating that observers' response criteria is such that they are equally likely to say a trial had or had not been altered. Conversely, in the unfamiliar condition, c was significantly greater than zero ($t(9)=4.24$, $p=0.002$), indicating that when observers are not familiar with a stimulus, they are predisposed to say that no alteration was present.

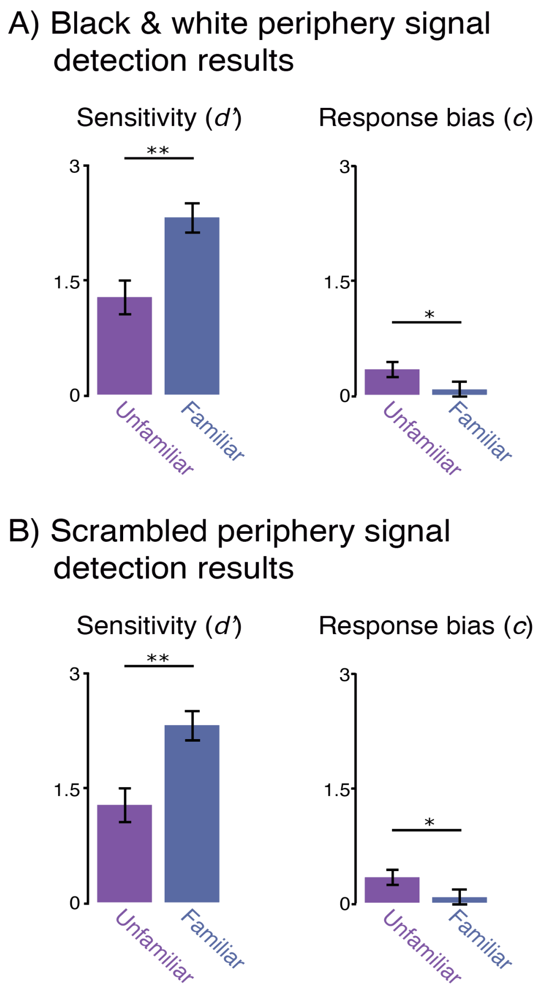

Figure 6. Experiment 2 results. A) Experiment 2a (black and white periphery). B) Experiment 2b (scrambled periphery. The y-axis represents *d'* (perceptual sensitivity) or *c* (response bias). Purple bars represent the unfamiliar condition and blue bars represent the familiar condition. Error bars represent the standard error of the mean. *$p<0.05$; **$p<0.01$.

Discussion

Overall, we found that when participants became familiar with a stimulus, their perceptual sensitivity increased while their response bias (or criterion) decreased. In the context of this experiment, when observers were not familiar with the target stimulus, they were relatively insensitive to noticing the peripheral alterations and were generally pre-disposed to saying that nothing had changed in the periphery. Once they were familiar with the stimulus, however, their sensitivity to noticing the alterations increased and they no longer had a bias one way or another towards saying whether something had changed in the periphery

Discussion

Across two experiments, we found that prior experience with a scene expands how much of that scene an observer is aware of. Specifically, observers were better at detecting alterations to the periphery when they were familiar with a stimulus compared to when they were not. These results are particularly noteworthy since no previous study examining the limits of perceptual awareness has manipulated stimulus familiarity in this manner. This is an important factor to consider when trying to understand the richness of perceptual awareness because in everyday life, observers are often highly familiar with their visual environment. Thus, in order to understand the capacity of perceptual awareness in naturalistic settings, researchers much consider how familiar an observer is with their surroundings. Indeed, these results suggest that by not considering this factor, researchers may be systematically underestimating the bandwidth of perceptual awareness.

What are the precise mechanisms that cause prior knowledge to expand perceptual experience? Although we cannot answer this question definitively, we believe that the mechanism most likely to cause this effect is visual attention (Chun & Wolfe, 2005; Carrasco, 2011). In situations when an unfamiliar stimulus is presented, the brief presentation of that target image (i.e., ~283ms) is not enough time for an observer to process the scene, identify informative items/locations, and direct their attention to those items/locations to determine if there is an alteration or not. However, when an observer is familiar with the scene, they can create a mental representation of the image and plan in advance where they will allocate their attention when the test image is presented (e.g., "The scene is of a large field of green grass that goes across my field of view, I'm going to attend far out to the periphery and see if I get any sense of green," or "There is a car on the edge of the periphery, I'm going to focus my attention there and see if its shape is scrambled on the test image"). This is likely a similar process that observers go through when reflecting upon and forming intuitions about the richness of their perceptual experiences. When considering how much information they are consciously aware of at a given moment, observers likely allocate their attention to specific parts of the visual world to investigate the bandwidth of their experience (e.g., "Hmm, some studies have said people don't see that much color in the periphery. I'm going to take a second and allocate my attention to the periphery and see how far I can see color. Yeah, that coffee cup on the end of my table is blue and I can see that it's blue in my periphery. I feel like I see more color in the periphery than prior results suggest."). Thus, we believe the allocation of attention potentially used in the familiar condition matches what observers often do when introspecting on the capacity of perceptual awareness.

It should be stressed that the process of allocating attention such that it expands the bandwidth of perceptual experience is a very different process from the one proposed by the filling-in hypothesis (Lau & Rosenthal, 2011; Chong et al., 2016; Otten et al., 2017; Tyler & Solomon, 2018; Toscani et al., 2021; Seth, 2021). Again, this view argues that prior knowledge of a scene leads to the visual system filling-in specific details that are not physically present (e.g., seeing an apple as red even though it is actually desaturated). By filling-in such details, this theory claims that observers actually *do* have a rich perceptual experience of the world around them, even if that experience is sometimes not veridical when experimenters alter the displays in one way or another (i.e., change blindness or inattentional blindness). Here, however, the pattern of results we found go in the opposite direction of what would be predicted from the filling-in hypothesis. While this by no means eliminates higher-level filling-in as playing an important role in perceptual awareness, it does lend serious challenge to the theory. Going forward, further empirical tests will need to be developed to continue exploring the upper limits of higher-level filling-in and its influence on the discrepancy between data and intuition.

Although these findings are inconsistent with the filling-in hypothesis, they fit nicely within a predictive coding framework (Rao & Ballard, 1999; Huang & Rao 2011; Clark 2013; Seth & Bayne, 2022). Under this view, the brain continuously creates and updates a model of the external world. To test the accuracy of this model, the brain compares how accurately the model predicted information that comes in through different sensory systems (e.g., visual, auditory, etc). The brain then calculates the difference between the predicted input and actual input and then updates the model to improve its accuracy. This process allows the brain to process information more efficiently and more easily adapt to the environment. The results reported here are consistent with this framework since familiarizing participants with stimuli allows the brain to generate a prediction about what an unaltered stimulus should look like. When the target image is shown, the brain has a clear template for comparing the internal prediction (i.e., the knowledge and memory of the unaltered stimulus) with the external sensory input (i.e., the target image). With such a prediction in mind, a predictive coding framework would suggest that participants would notice alterations more easily in the familiar condition. These findings lend support to the idea that what observers consciously perceive is, in part, the product of the brain's predictions (Hohwy & Seth, 2020).

Given that prior studies using change blindness and inattentional blindness did not ever use familiar stimuli, do the current results and framework described here significantly challenge or negate those prior findings (Rensink et al., 1997; Mack & Rock, 1997; O'Regan et al., 1999; Simons et al., 1999; Rensink, 2002; Most et al., 2005; Jensen et al., 2011)? Broadly speaking, we believe they do not. While it is possible that the rates at which observers would notice changes and unexpected objects in these paradigms would increase if the stimuli were highly familiar, we predict that those effect sizes would be relatively small and change and inattentional blindness would still routinely occur. However, this is merely a prediction and could be easily tested by presenting observers with images/displays, giving them time to familiarize themselves with those images/displays, and then immediately seeing how well they, for example, spot changes in a flicker paradigm. By directly comparing change blindness and inattentional rates with familiar and unfamiliar stimuli, we could indeed find that prior results using these paradigms have underestimated how much observers are aware of. Though, again, it should be stressed that just because familiarity may change the effect sizes of these paradigms, we do not believe those effects would by any means entirely disappear.

Finally, it is worth strongly emphasizing that we do not believe familiarity is the only factor related to the tension between empirical observation and personal introspection. Indeed, numerous other factors have been identified to help account for why observers feel as if they see more than the available data suggest. For example, one well-examined factor is known as subjective inflation, which refers to a set of findings showing that observers tend to exhibit overconfidence in both the periphery and unattended regions of space relative to central and attended regions even when performance is matched between them (Rahnev et al., 2011; Solovey et al., 2015; Odegaard et al., 2018; Knots et al., 2019). In addition, another potentially key factor is related to visual saccades. Specifically, although observers saccade roughly three times per second (if not more), they often fail to realize they are making any saccades at all and often drastically underestimate the number of saccades they make (O'Regan & Noë, 2001; Marti et al., 2015). Thus, people may often integrate across multiple eye movements and erroneously believe that all the information acquired during those movements was actually perceived simultaneously, giving the false impression of a richer phenomenology. Going forward, it will be important to understand the relative contributions of each of these different factors since together, they may close the gap between objective measurements and subjective experience.

References

Balas, B. & Sinha, P. (2007). “Filling-in” colour in natural scenes. *Visual Cognition*. 15, 765-778.

Beck, M. R., Levin, D. T., & Angelone, B. (2007). Change blindness blindness: Beliefs about the roles of intention and scene complexity in change detection. *Consciousness and Cognition*. *16*, 31–51.

Birch J. (1997) Efficiency of the Ishihara test for identifying red-green colour deficiency. *Opthalmic and Physiological Optics*,17, 403–408

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences,* 16, 567-575.

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

Carrasco, M. (2011) Visual Attention: The Past 25 Years. *Vision Research*, 51,1484-1525

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

Chong, E., A.M. Familiar & W.M. Shim. (2016) Reconstructing representations of dynamic visual objects in early visual cortex. *Proceedings of the National Academy of Sciences USA*. 113, 1453–1458.

Chun, M., & Wolfe, J. (2001). Visual Attention. In E. B. Goldstein (Ed.), *Blackwell Handbook of Perception* (pp. 2–335). Blackwell.

Clark, A. (2013) Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181-204.

Cohen, M.A. & Dennett, D.C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15, 358-364.

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, & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences,* 20, 324-335.

Cohen, M.A. & Rubenstein, J. How much color do we see in the blink of an eye? *Cognition*, 200, 104268 (2020).

Cohen, M.A., Botch, T.L., & Robertson, C.E. (2020) The limits of color awareness during active, real-world vision. *Proceedings of the National Academy of Sciences USA,* 117, 13821-13827

Cohen, M.A., Ostrand, C., Frontero, N., & Pham, P-N. Characterizing a snapshot of perceptual experience. *J. Exp. Psychol. Gen.*, 15, 1695-1709 (2021).

De Weerd, P., Smith, E., Greenberg, P. (2006) Effects of selective attention on perceptual filling-in. *Journal of Cognitive Neuroscience*, 18, 335–347.

Dehaene S (2014). *Consciousness and the brain: deciphering how the brain codes our thoughts*. New York, NY: Penguin Books.

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.

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

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.

Hohwy,J. & Seth, A. K. (2020) Predictive processing as a systematic basis for identifying the neural correlates of consciousness. Philosophy and the Mind Sciences. 1, 3.

Huang, Y., and Rao, R. (2011) Predictive Coding. *WIREs Cognitive Science*, 2, 580-593.

Jensen, M.S., Yao, R., Street, W.N., & Simons, D.J. (2011) Change blindness and inattentional blindness. *WIREs Cognitive Science,* 2, 529-546

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.

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.

Lamme, V.A.F. (2018) Challenges for theories of consciousness: seeing or knowing, the missing ingredient and how to deal with panpsychism. *Philosophical Transactions of the Royal Society B.* 373, 1755.

Lau H, Rosenthal D (2011) Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15, 365–373.

Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000) Change Blindness Blindness: The Metacognitive Error of Overestimating Change-detection Ability. *Visual Cognition, 7*, 397–412.

Loussouarn, A., Gabriel, D., & Proust, J. (2011). Exploring the informational sources of metaperception: The case of change blindness blindness. *Consciousness and Cognition*, 20, 1489–1501

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

MacMillan, N.A. & Creelman, D.C. (1991) *Detection theory: A user's guide.* Cambridge University Press: Cambridge.

Marti, S., Bayet, L., & Dehaene, S. (2015). Subjective report of eye fixations during serial search. *Consciousness and Cognition*, *33*, 1–15.

Matsumoto, M., Komatsu, H. (2005) Neural responses in the macaque V1 to bar stimuli with various lengths presented on the blind spot. *Journal of Neurophysiology*, 93, 2374–2387.

Maus, G.W. & Whitney, D. (2016) Motion-dependent filling-in of spatiotemporal information at the blind spot. *PLoS One*, 11, e0153896

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

Most, S. B., & Astur, R. S. (2007). Feature-based attentional set as a cause of traffic accidents. *Visual Cognition*, *15*(2), 125–132.

Odegaard, B., Chang, M. Y., Lau, H., & Cheung, S.-H. (2018). Inflation versus filling-in: Why we feel we see more than we actually do in peripheral vision. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *373*, 20170345.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*, 939–973; discussion 973-1031.

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

Otten, M., Pinto, Y., Paffen, C. L. E., Seth, A. K., & Kanai, R. (2017). The Uniformity Illusion. *Psychological Science*, 28, 56–68.

Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision,* 10, 437-442.

Pessoa, L., Thompson, E., Noe, A., (1998) Finding out about filling-in: a guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences*, 21, 723–748.

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

Ramachandran, V. S., & Gregory, R. L. (1991) Perceptual filling in of artificially induced scotomas in human vision. *Nature*, 350, 699–702.

Ramachandran, V.S., (1992) Filling in the blind spot. *Nature* 356,115.

Rao, R., & Ballard, D.H. (1999) Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2, 79-87.

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.

Rensink RA. 2002. Change detection. *Annual Review of Psychology*, 53: 245-277

Scholl, B. J., Simons, D. J., & Levin, D. T. (2004). "Change Blindness" Blindness: An Implicit Measure of a Metacognitive Error. In D. T. Levin (Ed.), *Thinking and seeing: Visual metacognition in adults and children.* (2004-18818-007; pp. 145–163). MIT Press.

Seth, A. *Being You*. London: Faber, 2021.

Seth, A. & Bayne, T. (2022) Theories of consciousness. *Nature Reviews Neuroscience,* 23, 439-452.

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

Solovey, G., Graney, G. G., & Lau, H. (2015). A decisional account of subjective inflation of visual perception at the periphery. *Attention, Perception, & Psychophysics*, 77, 258 – 271.

Toscani, M., Mamassian, P., & Valsecchi, M. (2021). Underconfidence in peripheral vision. *Journal of Vision*, 21, 2.

Tripathy, S.P., Levi, D.M., Ogmen, H., (1996) Two-dot alignment across the physiological blind spot. *Vision Research*, 36, 1585–1596.

Tyler, C. W., & Solomon, J. A. (2018). Does colour filling-in account for colour perception in natural images? i-*Perception*, 9, 1–10.

Ward, E.J. & Scholl, B.J. (2014) Inattentional blindness reflects limitations on perception, not memory: Evidence from repeated failures of awareness. *Psychonomic Bulletin and Review*, 22, 722-727.

Watson, A. B., & Pelli, D. G. (1983). Quest: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, *33*, 113–120.