Attribute amnesia or task-based interference?

A Thesis

Presented to

The Division of Philosophy, Religion, Psychology, and Linguistics

Reed College

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Arts

Caleb Kalisher

May 2016

Approved for the Division

(Psychology)

Michael Pitts

Acknowledgments

There are many people who have made it possible for me to be here writing this thesis.

To the SCALP lab, thank you for the constant source of camaraderie, even when I was the only person not doing an EEG study. Thank you, Shayan. I am not sure if we've ever met, but you gave me some code to base my stimulus off of when I had no idea how to program and without that I would have surely been lost. Chris, you were always there to help when I needed it eand you would frequently offer help or support even when I didn't realize I needed it. Thank you. Thank you, Enriqueta, for the advice throughout the process of this thesis. Michael, your guidance as my thesis advisor has been invaluable. Your enthusiasm for the research constantly reengaged my interests and helped me approach the problems that came up in a new light. I am incredibly grateful to all the time, and energy you gave me. Thank you.

To the friends I've made here at Reed who have helped keep me sane over these years, thank you. I would need far too many pages to properly express the gratitude that I owe you all. I would list your names, but I am not sure I could forgive myself if I left one off, so I trust that you will know who you are. Your friendship has truly meant the world to me, so thank you.

To every one of the Smooth Weasels for being the wonderful people that you are.

To my family for making Reed possible, for always reminding me to have fun, and for always being there for me, even when I forgot to call for far too long. Thank you.

Table of Contents

List of Figures

Abstract

Attention plays a crucial role for both perception and memory. While this is clear, there has been much debate over when and how inattention results in failures of perception and memory. A recent study (Chen & Wyble, 2015) introduced a paradigm that forced attention on an object by requiring a categorical discrimination (e.g., find the letter amongst numbers) prior to reporting the object's location. After performing the location-reporting task for a number of trials, participants were unexpectedly asked to report the identity of the object instead of its location (e.g., what letter was it?). The main finding was that participants were unable to report the identity of the object, even though they had presumably attended-to and perceived the object for the location task. These results were interpreted as inattentional amnesia for an object that had reached conscious awareness (Chen & Wyble, 2015). In an attempt to rule out alternative explanations, we modified this paradigm by increasing feature similarity between target and non-target objects to increase the attentional demands and we adjusted the way in which participants reported object location and identity to reduce potential interference during questioning. We hypothesized that the increased feature similarity would eliminate the inattentional amnesia effect by increasing the amount of attention given to the objects identity. Indeed, in our first experiment, participants reported the identity of the object well above chance when asked unexpectedly. In our second experiment, we returned to the original stimuli used by Chen & Wyble (2015), expecting to replicate their inattentional amnesia effect. However, we still found well above chance performance in reporting object identity on the surprise question, suggesting that reporting methods used by Chen & Wyble (2015) may have interfered with the participant's ability to report. It appears that determining whether an object was consciously perceived and then forgotten versus never perceived in the first place might critically depend on how perception and memory are probed.

Introduction:

1.1 Background

We like to believe that what we see is a complete representation of the world before us. This is why we have such idioms as “seeing is believing.” After all, if we didn’t see it how can we be sure it is there? Unfortunately, our assumption of perceptual perfection is incorrect, as research has shown that much of what goes on around us goes either unnoticed, or briefly noticed and forgotten.

Of course, hopefully, most of us accept that our perception and memory is not without its faults, as any number of day-to-day experiences can show errors in our senses. If our perception and memory were infallible we’d live in a world where eyewitnesses were always reliable and no one would ever misplace their keys.

In reality, we all miss quite a bit that goes on directly in front of us. One clear example of this is when you are so completely immersed in thought that what is happening around you fades to the background, despite your eyes being completely open. When you are finally shaken out of this state, perhaps by someone calling your name, you may apologetically say you were “spacing out.” This phenomenon is a common example of a lapse in our perception and could be considered a subset of perceptual lapses referred to as “sighted blindness” (Mack & Rock, 1998).

Examining the inner workings of our lapses in perception has been a focus of much research (Block, 2011; Bronfman, Brezis, Jacobson, & Usher, 2014; Cohen, Cavanagh, Chun, & Nakayama, 2012; Mack & Rock, 1998; Vandenbroucke, Fahrenfort, Sligte, & Lamme, 2014). By better understanding the strengths and weakness of the processes by which we perceive and remember the world around us, we can start adjusting how we approach the world to improve our absorption of information. Additionally, it will allow us to be more cognizant of when failures are likely to occur, and take that likelihood of memory or perception failure into account when asking for people’s recollection of objects or events.

While there is evidence for perception being possible without consciousness, i.e. "unconscious perception" (Pöppel, Held, & Frost, 1973) the relationship between attention and conscious perception is hotly debated. Some argue that before conscious perception can be achieved, we must pay at least some attention to the sensory input or the memory trace thereof (Cohen et al., 2012; Dehaene & Naccache, 2001; Mack, Erol, Clarke, & Bert, 2016; Mack & Rock, 1998). Others argue that attention and conscious perception are independent, with some types of conscious perception occurring in the absence of attention (Koch & Tsuchiya, 2012; Vandenbroucke et al., 2014).

Unfortunately, the nature of consciousness research is such that determining whether or not a subject consciously perceived a stimulus critically depends on subjective reporting. Determining the cause of the absence (or inaccuracy) of a subjective report is often difficult and suspect to multiple interpretations. For example, if someone reports not seeing an object, it could be due to a failure to consciously perceive the object in the first place, or to a failure to remember the perception of that object by the time they are asked to report it. In either case, the report of absence is the same, despite conscious perception being present in one case and not the other (Block, 2011; Bronfman et al., 2014; Keizer, Hommel, & Lamme, 2015; Koch & Tsuchiya, 2012; Vandenbroucke et al., 2014).

This division of explanations has resulted in much debate within the vision research community as to the source of our apparent "blindness." Even if both camps agree that a primary cause is inattention, some have proposed that perception is disrupted by attention (Mack & Rock, 1998), while others have argued that inattention leads to an immediate loss of memory, or amnesia, despite completion of perception (Wolfe, 1999). To address this debate, the present research focuses on a newly developed experimental paradigm (Chen & Wyble, 2015) in order to parse out the possible explanations for an absence of report.

1.2 Attention and Memory

It should come as no surprise that attention plays a critical role in our ability to remember. The dependence of memory on attention has been demonstrated in numerous different ways (Halin, Marsh, & Sörqvist, 2015; Lavie, Hirst, de Fockert, & Viding, 2004; Ruz, Worden, Tudela, & McCandliss, 2005). It has been found that individuals asked to complete a letter recall task while simultaneously listening to a story show a reduction of ability to report on information regarding the story. This is attributed to the target task demanding enough of working memory's resources to inhibit the subject's attentional capacity for the distractor story (Halin et al., 2015). A study by Palmer, Boston, and Moore (2015) had people identifying features of an object. The greater the number of features on the object, the fewer total number of features subjects were able to remember. This reduction in memory performance was ascribed to the increase of attentional demand required by objects with more features (Palmer et al., 2015).

As demonstrated by the Halin & Marsh (2015) and Palmer et al. (2015) studies, the diminished ability to remember is at least partially determined by the resources being used by working memory, also known as "cognitive load". The idea of cognitive load is that attention is a limited quantity mental resource. This pool of attention is mediated by limitations of both our working memory (Lavie et al., 2004) and physiological needs (Sünram-Lea, Foster, Durlach, & Perez, 2002).

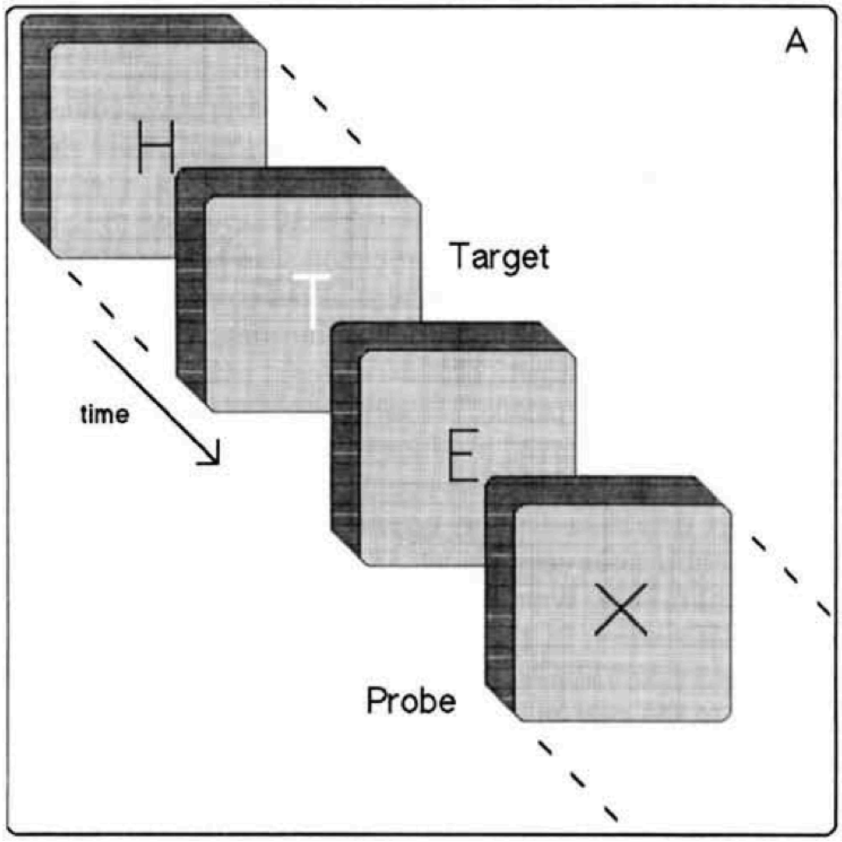

Figure 1.1 RSVP Task for the Attentional Blink
Above is a visual representation of the RSVP task from Raymond et al. 1992. Subjects were tasked with attending to the white printed letters. They were then asked to respond if the black X (the probe) appeared in the subsequent stream of letters.

One phenomenon that comes out of the limitations of attention on memory is the attentional blink (Shapiro & Luck, 1999). An attentional blink occurs when short-term memory is overloaded by attentional tasks, causing a brief period of time (200 – 400 ms) where the transfer of information to memory appears to be disrupted, while visual perception is still intact (Luck, Vogel, & Shapiro, 1996).

Attentional blinks are generally created in the lab using a paradigm known as a rapid serial visual presentation (RSVP). The simplest form of the RSVP is done by presenting a subject with a stream of stimuli on a screen. While the images are being presented, they are told to identify a certain stimulus that is somehow differentiated (such as a different color). The attentional blink is found when subjects are asked to identify a second stimulus in the stream (Figure 1.1). Despite subjects being able to successfully report on either of the two stimuli by themselves, when they are asked to identify both stimuli there is a window of time where individuals are unable to report the second stimulus (Raymond, Shapiro & Arnell, 1992).

With attention being a limited commodity, if we are focusing on a primary task or stimulus, the likelihood we will be able to see or remember seeing anything else diminishes as a function of the amount of attention devoted to the primary task or stimulus (Lavie et al., 2004; Mack & Rock, 1998). Cognitive load results in two additional phenomena that will be discussed in greater detail over the next few sections: inattentional blindness and inattentional amnesia.

1.3 Inattentional Blindness

As attention to one object or event influences what else we report seeing, it is important to examine what occurs when absolutely zero attention is allocated to a stimulus. Inattentional blindness occurs when we fail to perceive stimuli that are outside the focus of our attention. The discovery of inattentional blindness led Mack and Rock (1998) to formulate the hypothesis that "there is no conscious perception without attention." What this means, is that for us to be aware of an object or event we must, at least partially, be attending to that object or event. If we are not paying attention, according to Mack & Rock's (1998) hypothesis, we cannot perceive that object or event and are therefore "blind" to it.

It turns out that we are very easily made inattentionally blind. It is one of the reasons why magicians can get away with so many shenanigans (Macknik, Martinez-Conde, & Blakeslee, 2010). Magicians call it sleight of hand, but we know the truth, it's really "sleight of mind." The typical design of studies that create inattentional blindness is to have subjects focus on one task, and occasionally something else unexpected will happen (they are not initially told about the unexpected object or event). There are simple studies, such as one done by Arien Mack and Irvin Rock, where people fail to see shapes appearing on a screen when they are asked to focus on a cross with changing arm lengths in the middle of the screen (See Figure 1.2) (Mack & Rock 1998).

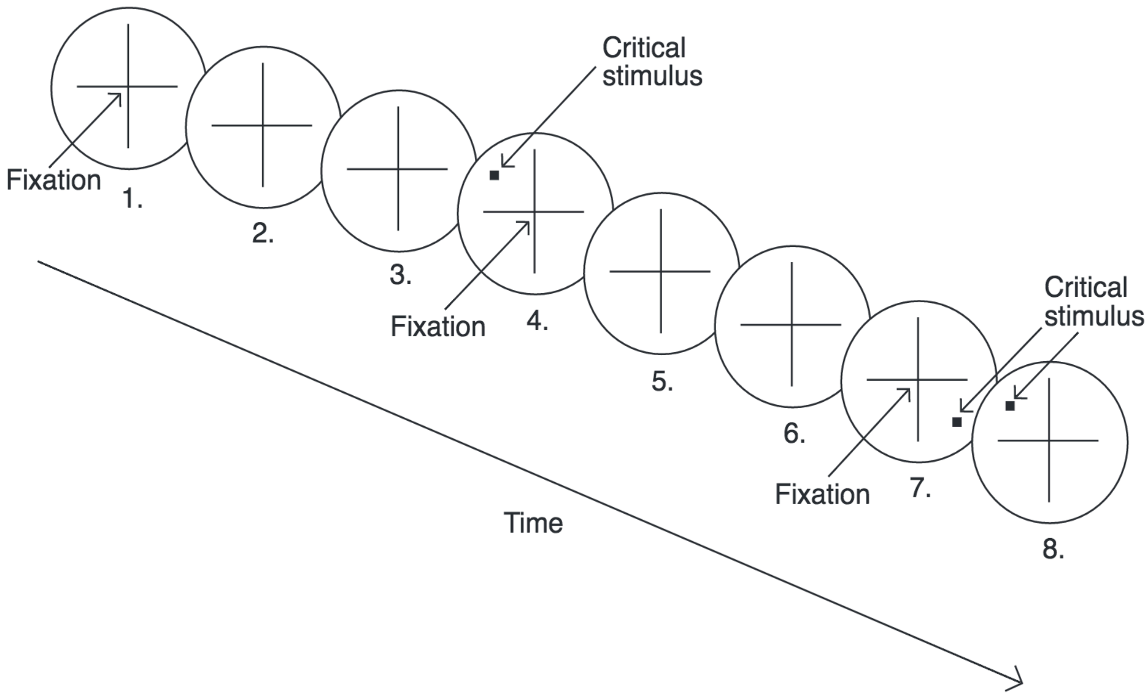

Figure 1.2 Visual representation of an inattentional study
Subjects are tasked to observe changes in the length of the fixation cross. In doing so they are made unaware of the "critical stimulus" appearing on the screen (Mack & Rock, 1998)

There are also many studies in which researchers have tried to push the limits of inattentional blindness. In the most famous inattentional blindness study, Daniel Simons and Christopher Chabris (1999) were able to make subjects completely blind to a person in a gorilla suit walking across the screen by simply having subjects focus on counting how many times players in black or white shirts passed a basketball.

While less dramatic than an ape trouncing across a screen, a recent study by Emily Ward and Brian Scholl (2015) provides excellent evidence for the pervasiveness of inattentional blindness. Their study had subjects in a task that involved tracking the movement of black or white letters on a screen to count how many times they crossed the midline. While they were doing this, a blue cross or a red E would move across the screen (for 10 seconds). For a visual representation of the task, see Digure 1.3. Many subjects (30%) did not notice the blue cross or red E. However, after this unexpected event occurred, subjects were asked to report if they saw any "unexpected events" in the following trials. Despite this prompting to be on the lookout for the unexpected, a fair number of subjects (25%) still did not report anything out of the ordinary, suggesting that telling someone to "expect the unexpected" does not substantially improve one's ability to notice things. According to this study, in order to notice a stimulus, you either need to know exactly what to look for in advance, or happen to have you attention in the right place at the right time.

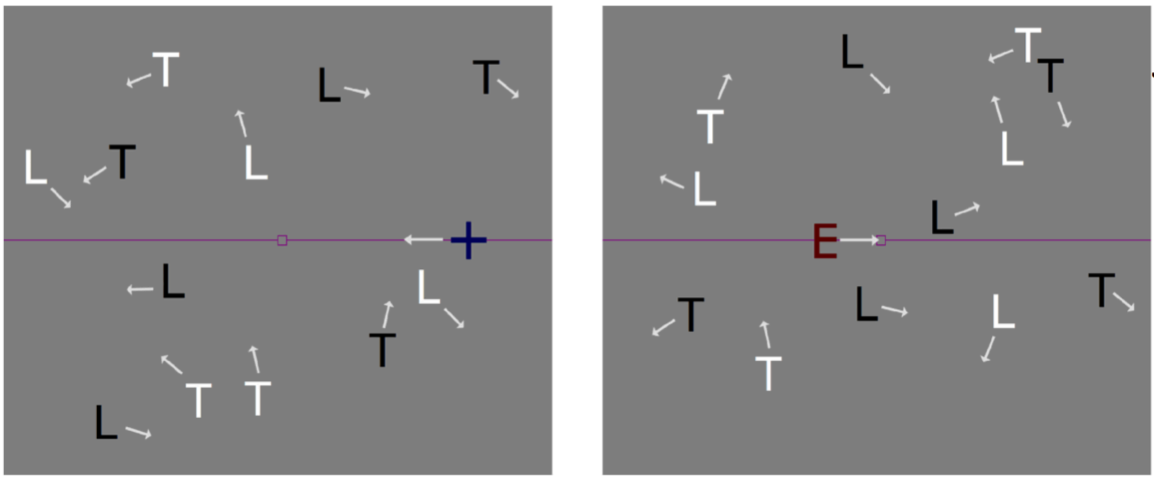

Figure 1.3 Examples of Unexpected Events from Ward & Scholl study (2015)
Above are two examples of the stimulus that subjects were watching (the white and black "T"s and "L"s) and two unexpected stimuli that could occur (the blue cross and the red "E").

1.4 Objects and Features

In the previous examples, we have been considering stimuli as singular objects or events. This, of course, is not entirely accurate. There are in fact many elements, or features, that compose these stimuli. In Simons and Chabris (1999) study the target task involves multiple events between multiple people. These people are wearing different colored shirts (black or white) and are passing around a dark orange basketball, thus color and motion tracking are essential elements of this task. The gorilla is not just an image that pops up on the screen. It is a woman in a gorilla suit walking across the screen (and stopping to pound her chest halfway through). It is the motion, shape, color, texture, size, and arrangement of facial features that combine to give us the perception of the gorilla.

Even in simple stimuli there are many different features that create what we perceive. A small square was used as a critical stimulus in some of Mack & Rock's experiments (Figure 1.2), which can be defined by its shape, size, position, and whether it is filled in or textured. The letters used in Ward & Scholl's (2015) experiment are composites of curved and/or straight lines of certain lengths that form intersections or end-stops, along with whatever color is used to fill in the letter. For example a "T" is generally composed of two lines, with a long vertical line meeting a slightly shorter horizontal line at one of its ends. While it may seem silly to break down simple stimuli in such a manner, there has been evidence that shows limitations in memory in regards to the number of objects, and the number of features those objects contain (Palmer et al., 2015; Persuh, Gomez, Bauer, & Melara, 2014).

There have been primarily two contending hypothesis for how we process objects and features for a number of years: the discrete-object limit hypothesis and the feature-precision limit hypothesis (Palmer et al., 2015). The discrete-object limit hypotheses claims that, memory being discrete, we can only remember a given number of objects. Therefore, we will store all the features that compose those objects (Luck & Vogel, 1997). On the other hand, the feature-precision limit hypothesis interprets the discrete nature of memory as a result of a fixed total amount of information, with individual features being the base measurement and therefore the limiting factor (Palmer et al. 2015). In other words, we might be able to remember six objects but with less precision (perhaps by remembering fewer of the objects features) or we could remember three objects with high precision (including all features).

While these hypotheses are making claims toward the memory capabilities in regard to objects and features, we must also examine the perceptual portion of information processing (necessarily before memory can occur), for a mutualistic understanding of the processes of perception and memory. A study by Persuh et al. (2014) found evidence for feature-based inattentional blindness. What that means is that people were able to perceive and report seeing part of a stimulus but were blind to another feature of that same stimulus.

There has also been evidence for the level of attention needed to perceive objects and different types of features. Rock and Gutman (1981) found that shape related features for a "novel form" (see Figure 1.4 for example) required attention to perceive, while certain features, such as color or size of the object, could be perceived "for free" even if the subject was attending to a different feature of that object.

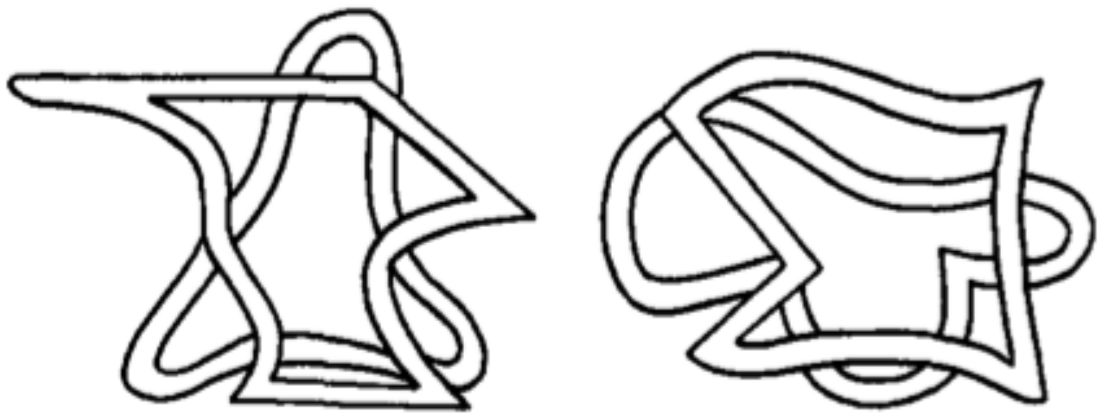

Figure 1.4 Example of "Novel Form" Objects
Above are two example of stimuli presented to subjects in the Rock and Gutman (1981) study. They would be filled in with a color.

1.5 Inattentional Amnesia

One of the problems with inattentional blindness paradigms is that the tests used to assess perception must be administered after the fact, because the entire effect is centered on the unexpected nature of the critical stimuli. Because the questions about what was perceived are necessarily delayed, memory must also be involved. This memory requirement creates a question that must be asked: What if we momentarily perceived the critical unexpected stimulus, but failed to retain it in memory long enough for us to report it after the fact?

This question has given rise to the inattentional amnesia hypothesis. The amnesia hypothesis claims that in standard circumstances we consciously perceive most of what is happening in our field of vision, however information that we are not attending is much less likely to be held in working memory or consolidated into long-term memory (Wolfe, 1999). There have been findings that suggest that our ability to transfer visual stimuli into memory is inhibited by high attentional loads (Ruz et al., 2005), such as other distracting stimuli or difficult tasks.

There have been a number of experiments that intuitively make more sense with inattentional amnesia as an explanation rather than inattentional blindness. One study by Simons and Levin (1997) had someone ask a random person for directions. After engaging in conversation for a while, construction workers would walk between the two people carrying a door. The person asking for directions would be switched out with a completely different person, and 50% of the time the subject would fail to realize this change. As engaging in conversation would presumably require someone to attend to and consciously perceive the other individual, the idea that they never consciously perceived the person in the first place is a little harder to accept. The amnesia view, on the other hand, posits that the degree of attention allocated to the person was not great enough to create a concrete memory of the person, hence the failure to report the swap.

In another study, subjects watched a video in which a woman poured Diet Pepsi from a bottle. The bottle could be seen clearly for approximately 6.5 seconds, until the camera panned toward an approaching individual. When the camera returned, a cardboard box had replaced the bottle. When subjects were asked if anything unusual occurred, no one reported the switch, despite multiple subjects reporting on the nature of soda, or how it was poured (Simons, 1996).

The difficulty that quickly arises when trying to assess the inattentional amnesia hypothesis is that the end product is virtually indistinguishable from that of inattentional blindness. With either explanation, people are left unable to report seeing an object or event or change that was presented to them. To better test these two alternatives (blindness vs. amnesia), it may be possible to create a stimulus and task that would require a subject to consciously perceive an object, while still rendering subjects unable to report seeing the object. Such a set-up would eliminate the explanation of inattentional blindness (i.e. they couldn't do the task without consciously perceiving the stimulus) while providing a clear example of inattentional amnesia.

1.6 Chen and Wyble

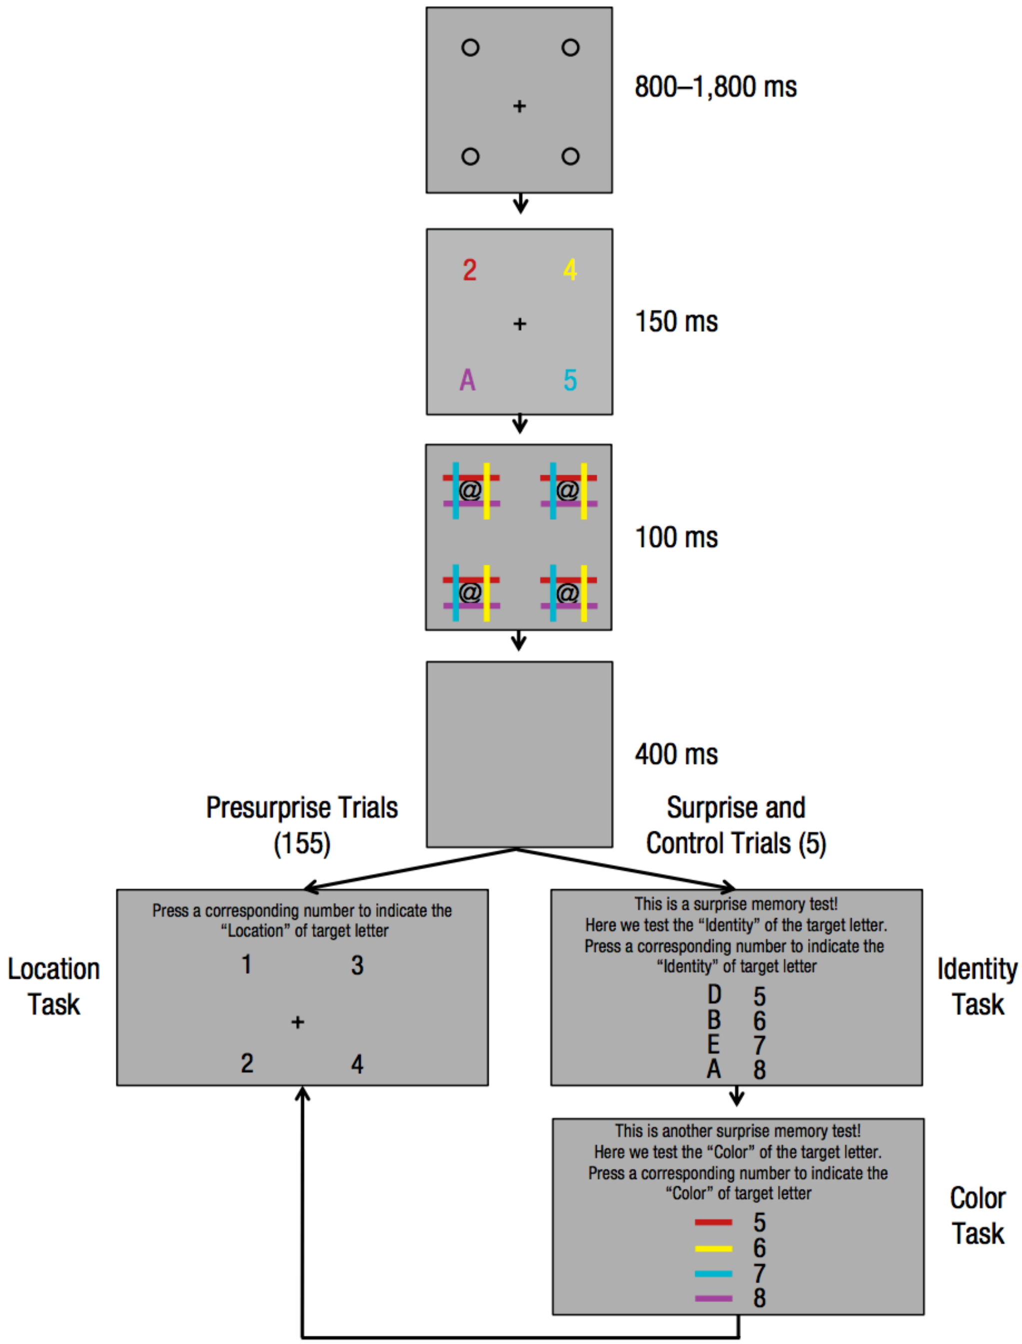

Figure 1.5 Chen and Wyble's experimental design

This is the stimulus used by Chen and Wyble (2015). To the right of each screen is the duration of time each screen was presented in milliseconds. The first screen shown to subjects was a placeholder screen to act as buffer and to show possible location of letters and numbers. The 2nd screen shows the target stimulus (the letter) amongst numbers. The 3rd screen was a mask. The 4th screen was a blank buffer before the subject was asked either the pre-surprise question about the location of the letter (left), or the surprise questions about the identity of the letter and its color (right). After the surprise question, they were also asked the question about the letter's location. This surprise trial was repeated a few times to see if subjects could report the identity and color of the letter if they knew to attend to it ahead of time (dual-task control trials).

Chen and Wyble attempted to create a paradigm that controlled for inattentional blindness while testing for inattentional amnesia. Their study involved a stimulus and task that would require subjects to consciously observe the target, but fail to report some of the features of that object.

To do this they created a task (see Figure 1.5) in which subjects had to find a single letter amidst three numbers, and report the location of this letter. On each trial, one out of four possible letters (A, B, D, E) served as the target, and three out of four possible numbers served as the non-targets (2,3,4,5). The four stimuli were randomly positioned at the corners of an imaginary square centered around fixation and were drawn in different colors (red, blue, yellow, magenta). On the pre-surprise trials, subjects were asked to indicate the location of the letter. After 155 pre-surprise trials, the location-task was interrupted and two surprise questions were asked immediately after the letter/number display, one that asked for the identity of the letter, and another that asked for the color of the letter. Presumably, for a subject to identify the location of the letter, they must have identified the letter. Chen & Wyble argued that this location task required consciously attending to the entire target stimulus. The surprising result was that, while subjects could identify the location of letters correctly (89% accuracy), they were at chance when asked to identify of the letter (25%) or the color of the letter (30%).

Importantly, in a follow-up experiment with a new group of naïve subjects, Chen & Wyble (2015) reduced the number of pre-surprise trials to 11 (instead of the previous 155), and found the same results. They also repeated the study giving subjects a 100 ms increase in time to perceive the letter (total of 250 ms) and removed the mask in an attempt to increase attention to and awareness of the letters. In another version, they ran the study having people identify odd numbers (13, 15, 17, 19) amongst even numbers (12, 14, 16, 18) and vice versa. In all of these cases they found that people were at chance in reporting the identity of the target, despite good performance on the location task.

They conclude that, as the pre-surprise trial task of identifying a letter or number's location necessitated conscious perception of that stimulus, the inability to report either identity or color was due to a memory failure, not a perceptual failure. This interpretation implies that attention was insufficient for creation of a memory, and that the filtering of relevant and non-relevant information occurs after conscious perception.

Attribute Amnesia

Chen and Wyble explain their findings through a subtype of inattentional amnesia they call "attribute amnesia". Attribute amnesia would occur when a subject consciously perceives an object, including all of its attributes, but a second or two later, is unable to report one or more of the object's attributes. In their 2015 study, attribute amnesia is said to have occurred when subjects could report the location of the target letter with great accuracy, but were unable to report the identity or color of that letter.

This explanation follows the logic of the feature-precision hypothesis (Palmer et al., 2015), which allows for attribute amnesia, as the attention required for the location task would reduce the memory capacity for task-irrelevant features.

Additionally, there have been some findings that suggest that there is an initial pre-attentive perception of unbound visual features, followed by an attention-dependent "binding" of features to the object as a whole at a particular location (Treisman & Gelade, 1980; Wolfe & Bennett, 1997). If this is true, it would demand that, in Chen & Wyble's experiments, subjects perceived the bound features (letter shape and color) because they had to attend to them to perform the location task. By taking these two theories together, we have perception of object in a certain location with a limited memory for features, or attributes of those objects.

Criticism

Chen and Wyble make an admiral attempt to design an experimental paradigm that allows a parsing out of the difference between inattentional blindness and inattentional amnesia. While this is clearly a step in the right direction, a number of potential concerns must first be addressed.

One of the clearest counter explanations to their finding is that subjects may have adopted a strategy in which they only attended, and thus only perceived, a high-level category (e.g. "letterness" or "numberness") in order to carry-out the location task. This could manifest in various ways in their paradigm. Subjects may have looked for the location of the symbol that fits into the broad category of "letter" while not actually attending to the features of the letters and numbers. This strategy could be facilitated by our extreme familiarity with letters and numbers as distinct categories. There is evidence that supports perception of categories occurs before the perception their sub-categories (King & Dehaene, 2014), which would provide an alternative explanation of identifying the "letterness" category of the target stimulus before the subject has to identify the letter. If this is the case, "letterness" would be sufficient for the location task, and therefore, conscious perception of the identity of the letter is not mandatory. Another possible strategy using this top-down model of perception is that subjects may have been able to identify "numberness" of the distractor stimuli, in order to eliminate possible locations of the "not number." This process-of-elimination approach could be achieved without ever attending to the letter itself. After all, the stimuli were flashed very briefly, in the visual periphery, and there were only three locations to rule out in order to accurately determine the target location.

Chen and Wyble attempt to dismiss this high-level categorization problem with their second experiment, in which the task was to find an odd number amongst even numbers, or vice versa. The problem here is that odd and even numbers are still distinct categories. This of course implies that it may not be possible to use their task to address this problem, as no matter what stimuli are used, one will always have to identify an object of one category amongst objects of another category. What their second experiment does manipulate is discrimination difficulty, and by doing so it may increase the amount of attention required to locate the target object. However, the features that allow us to differentiate between a letter and a number, or between an odd number and even number, may still be distinct enough to facilitate high-level categorization.

Another key concern with Chen and Wyble's experiment is in the methods that they employed to collect responses, both for the location task and the surprise identification questions. In the location task, the four possible response options (upper-left, upper-right, lower-left, lower-right) were listed on the screen as numbers (1, 2 ,3, 4) that appeared at the each location. Subjects were tasked with pressing the corresponding number keys on a keyboard to indicate the location of the target stimuli. In the surprise questions, the four options were listed beside four additional numbers (5, 6, 7, 8) and subjects were tasked with typing the number that corresponded to the letter that had just been presented. In other words, the reporting system required subjects to convert location and identity (and color) to a number system, even when numbers were a core part of the stimulus to begin with. This requirement to convert a perceived feature to an arbitrary numeric code for response purposes may have interfered with memory, causing the apparent attribute amnesia effect.

1.7 Present Study

While Chen and Wyble's research might suggest a new phenomenon closely related to inattentional amnesia ("attribute amnesia"), its shortcomings must now be addressed. Our research tests several potential confounds in Chen & Wyble's studies in an effort to get us closer to distinguishing between inattentional amnesia versus inattentional blindness.

Because of the unique requirements for inattentional amnesia and attribute amnesia to exist, it is extremely difficult to test. The task must require the participant to consciously perceive the target as a whole, while only being asked about one feature within the target. In this way, the use of categories seems unavoidable. The trick then becomes finding a way that demands greater attention of the participant to subcategorize the target from the distractors.

Our first study attempts to make the features of the categories more similar in hopes of reducing ease of differentiation, and therefore increasing the amount of attention required.

Rationale

This study contains multiple changes to Chen and Wyble's (2015) attribute amnesia paradigm.

The first change was to make the point when the surprise question is asked contingent upon how well people are performing in the location task. In Chen & Wyble's experiments, a set number of trials was administered before asking the surprise question (in some cases 155 pre-surprise trials, in other cases 11), without any contingencies based on performance. It seemed critical to ensure that when a participant is unable to answer the identity question that it is not because they couldn't locate the target stimuli in the first place (e.g., we wanted to control for brief attentional lapses, individual differences in task performance, and shifts in sustained effort).

The second change was to simply alter how the participants report location and letter identities. We were concerned that having participants answer questions involving numbers and letters with number input commands (e.g., 2 = Top Left and, 5 = B) may be causing interference in working memory thereby impacting the main results. Instead of requiring participants to convert their location and letter identity responses to numbers, we simply asked them to click with the mouse cursor on the location of the screen in which the letter appeared and to type the actual letter into the keyboard.

Finally, in our first study we increased the similarity of features between the letters and numbers by carefully choosing a font style and letters and numbers that shared many component features. This change in the physical properties of the letters and numbers was designed to require participants to pay more attention to the letter identity in order to discriminate letters from numbers, thereby ensuring that they were aware of the letters prior to potentially forgetting them when asked the surprise question. In other words, this change made it more difficult for participants to perceive the "letterness" and "numberness" category without perceiving the specific letters and numbers.

Predictions

The first possible result of Experiment 1 is that the increase of feature similarity will result in people being able to identify the letter in the surprise question (i.e., the attribute amnesia effect will disappear). If this is the case, Chen & Wyble's (2015) findings might be better explained by inattentional blindness to letter identity.

The second possibility is that we find attribute amnesia effects similar to Chen and Wyble. This would mean that despite the increase in feature similarity between letters and numbers, people are still only able to report letter identity at around chance levels. This result would lend support to the inattentional amnesia hypothesis, assuming that our changes to the physical stimuli succeeded in requiring participants to maximally attend to letter identity in order to complete the location task.

Present Study

We examined the hypothesis that individuals are rapidly forgetting the identity of an object, even when they are required to attend to the attributes of the object. Participants performed a simple visual memory task nearly identical to the Chen and Wyble (2015) study, asking them to look for a location of a single letter stimulus amongst number stimuli distracters. After proficiency in the task was demonstrated, participants were asked to identify the letter on a surprise trial. Several control trials followed the surprise trial, in which both location and identity were reported. The goal of this task is to require participants to consciously attend to the target letter, which means that if they cannot report the identity of the letter this failure to report is likely due to amnesia rather than perceptual blindness.

2.1 Experiment 1

In Experiment 1, the font, letters, and numbers where changed to decrease ease of participant's ability to discriminate between the letters and numbers. Each letter was matched with a number that had familiar features in the font that we used. The goal of this change is to reduce the possibility that a categorization effect is allowing participants determine the location of the letter without having to consciously identify the letter.

Methods: Experiment 1

Participants

Participants were 30 Reed college students, all over the age of eighteen. All the participants gave their informed consent prior to the start of the experiment. Participants were compensated for their time with lottery ticket that gave them the opportunity to win $50 in a departmental drawing. Experimental approval was obtained from the Reed Institutional Review Board before data collection began.

Apparatus

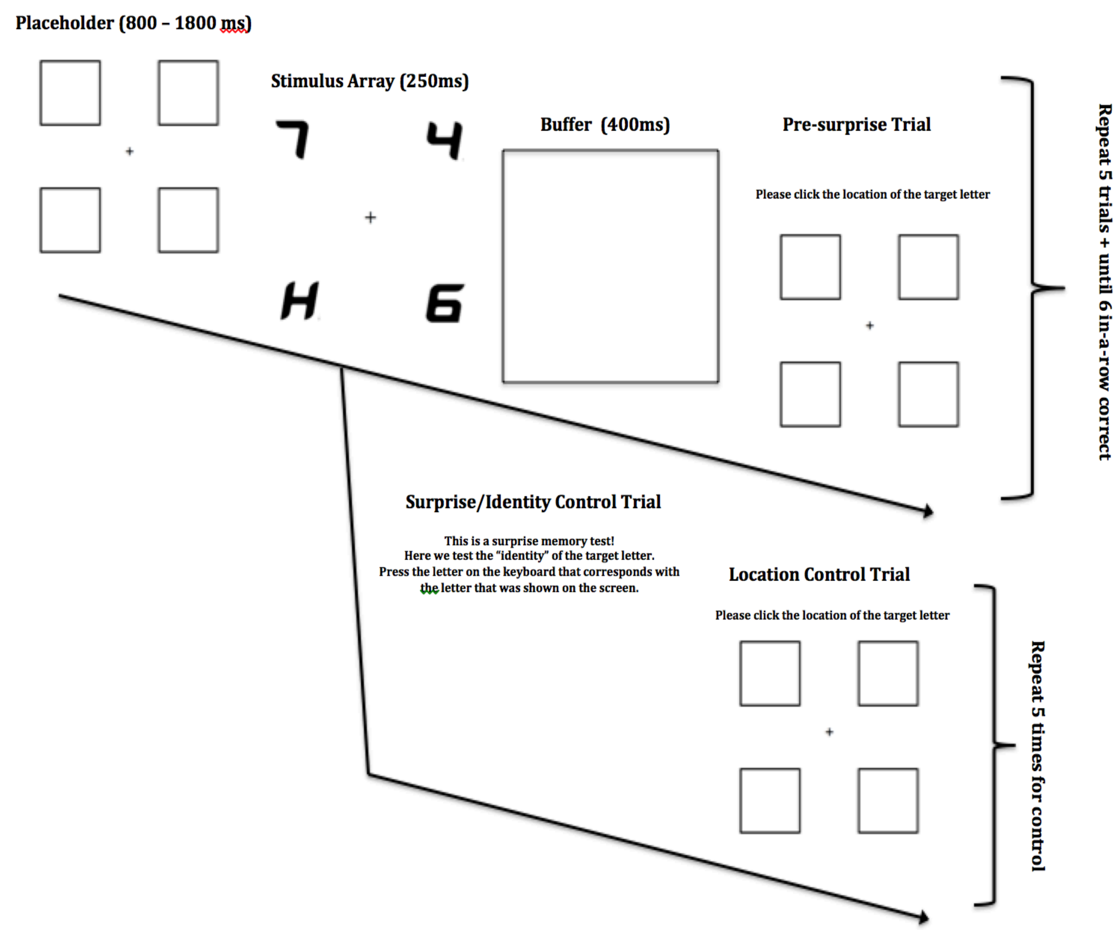

Figure 2.1 Example sequence of stimuli
The order of stimulus presentation is from left to right. First there is a placeholder image for 800-1800ms, followed by the stimulus array for 250ms. There is a 400ms blank buffer before the pre-surprise question. After participants reach the contingency requirement (6 pre-surprise trials in-a-row with correct location reports), they are asked the surprise trial (bottom left) instead of the pre-surprise trial. The surprise question is followed by a location question identical to the pre-surprise question. This sequence of asking for letter identity and then location is repeated five times as a dual-task control.

The stimuli were presented on a 61 cm computer monitor. Participants were seated so their eyes were approximately 57 cm away from the screen. The participants' responses were entered via either mouse (for locational tasks) or keyboard (for identity tasks). The stimuli was created and presented through Presentation software (Neurobehavioral Systems, San Francisco CA).

As can be seen in Figure 2.1, each trial started with a screen with a fixation cross in the center and four black placeholder squares on a grey background. Each square was located in the corner of an invisible square at the center of the screen that was 6.5° x 6.5° visual angle size. This placeholder screen was presented for a variable interval of 800 – 1,800 ms before the stimulus array was presented.

The stimulus array consisted of three numbers and one letter in the "Bloodwax" font, each character located in the same position as the placeholder squares in the previous screen (Figure 2.1). The stimulus array is presented for 250 ms.

The numbers displayed randomly with the option of 4, 3, 7 and 6 (Figure 2.2). The letter displayed randomly from the options of H, E, L, and P (Figure 2.2). Each character on the screen was approximately 0.86° x 0.86° of visual angle size. The target letter and distracter numbers were selected completely at random on each trial, and positioned randomly in the four possible locations.

Figure 2.2 Example of Possible Characters
Above are the letters and numbers in the "Bloodwax" font used for Experiment 1. Each possible letter is shown above the number that it was matched with for feature similarity.

After the stimulus array, a 400 ms blank screen was presented as a buffer before the pre-surprise question. The pre-surprise questions present the placeholder screen, along with a 10 pixel square cursor in the center of screen. Participants were asked to move the cursor to the square where the letter was located and click to register their response. After the participant answered the pre-surprise question, a new trial started from the beginning. The participant completed this type of trial 5 times (as a warm-up) before the contingency counter began. Once they answered 6 additional pre-surprise questions in-a-row correctly, they were asked the surprise question (thus, a minimum of 11 pre-surprise trials were presented to maintain consistency with Chen & Wyble's (2015) Experiment 4). The surprise question tasked the participant to enter the letter on the keyboard that they just saw. They were then asked to report of the letter. The identification followed by location question was then repeated for 4 additional (dual-task control) trials.

Results: Experiment 1

Participants showed high accuracy on the pre-surprise location task. They were able to correctly identify the location of the letter 94% of the time, taking on average 7.96 pre-surprise trials to reach the contingency requirements for being asked the surprise question (6 in-a-row correct). Thus, an average of 13 pre-surprise trials were administered (minimum = 11: maximum = 21). This indicates that participants were able to quickly learn the task and locate the letter with ease.

When asked to identify the letter on the surprise trial, participants could do so with high accuracy. They were able to correctly identify the letter 83% of the time. Additionally, they were still able to identify the location of the letter above chance, at 76% accuracy, even though the surprise question about identity was likely to interfere with the location task (see Figure 2.3).

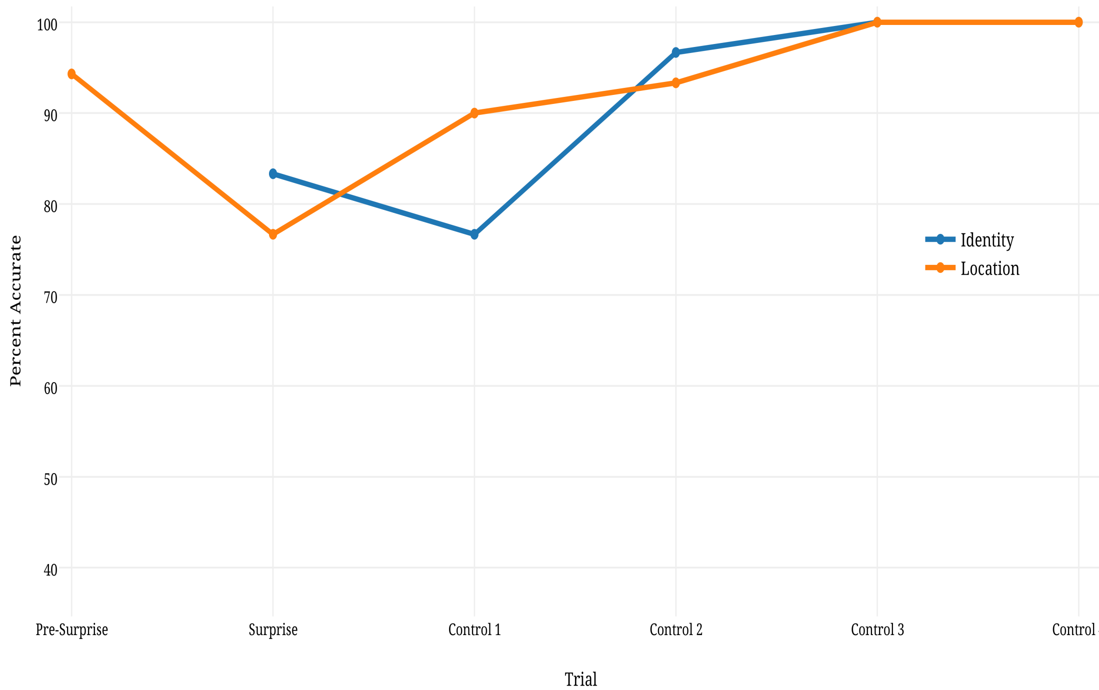

Figure 2.3 Experiment 1 Accuracy Results

Identity and location accuracy compared across the pre-surprise, surprise, and control trials. Pre-surprise accuracy is an average from all of the contingency location trials. The slightly lower-levels of accuracy on the surprise and 1st control trials are likely due to the task-switch costs (single-task to dual task). Despite these minor decreases in accuracy, participants were well above chance on the surprise accuracy trial, thus failing to exhibit attribute amnesia.

Participants also showed positive trends of accuracy reporting during the control trials. There was an initial drop of letter identification accuracy from 83% in the surprise trial to 76% in the first control trial, however this was not significant, $\chi^2(1, N = 30) = 0.104$, $p = 0.74$. Accuracy then increased to 96% on control trial 2, and reached and maintained 100% accuracy on control trial 3 and 4. Location accuracy returned to near pre-surprise levels of accuracy (94%) on the control trials, with control trial 1 accuracy being 90%, control trial 2 accuracy being 93%, and control trial 3 and 4 reaching 100% accuracy.

Discussion: Experiment 1

The data from Experiment One show no hints of the attribute amnesia effect found by Chen and Wyble study (2015). This is evidenced by the participants being able to report the identity of the letter on the surprise-trial with great accuracy.

This finding leaves two primary explanations for Chen and Wyble's findings. First, in support of a top-down model of perception, we can perceive the "letterness" of a letter before we perceive which letter is present. This would allow participants in the Chen and Wyble study to have skipped identifying the letter to find the location of the "letter-like" object. In the current study, the categorization of "letterness" was made more difficult by increasing feature similarity of the numbers and letters.

The second possible explanation for the discrepancy between the current findings and those of Chen & Wyble pertains to the way they were asking the questions. By asking their participants to translate position and identity into arbitrary numbers they may have been increasing the cognitive load and interfering with the memory of a just=perceived individual letter.

To determine which of these explanations best explains the results, we ran experiment two.

2.2 Experiment 2

As Experiment 1 showed that participants were able to report the identity of the letter in the surprise trial, we now sought to re-create the attribute amnesia effect found by Chen and Wyble. For this study we returned to the same stimuli used by Chen and Wyble retaining only our changes to the response input methods (click the location, enter the letter). We predicted that the dissimilarity of features between letters and numbers will allow the "letterness" strategy to be used to find the location resulting in inattentional blindness to the actual identity of the letter.

Methods: Experiment 2

Participants

Participants were 30 Reed college students, none of which participated in Experiment 1. See section 2.1.1 for more detail.

Apparatus

See section 2.1.1.

Stimuli

The stimulus was identical to the previous study (see 2.1.1) with the exception of the letters and numbers and their font. For this study the stimuli were the same letters and numbers from Chen and Wyble's (2015) study (A, B, D, E and 1, 2, 3, 4). We also used the Times New Roman Font instead of the "Bloodwax" font used in Experiment 1.

Results: Experiment 2

Participants were able to correctly locate the letters 100% of the time (6 times in a row) in the contingency pre-trials. Despite the flawless pre-surprise accuracy for location, 10 out of 30 participants were unable to identify the letter on the surprise question. This meant that letter identity for the surprise task was at 67% accuracy, still well above chance (25%). There was also a 10% drop in accuracy for location accuracy in the surprise task versus the perfect pre-surprise accuracy.

The post-surprise trials showed rapid improvement in letter identification accuracy, and a return the pre-surprise location accuracy. Both of the measures showed a 99% accuracy across the four post-surprise trials. Unlike experiment 1, the improvement from the surprise trial (67%) to the control 1 trial (99%) was significant, $\chi^2(1, N = 30) = 9.72$, $p = 0.0018$.

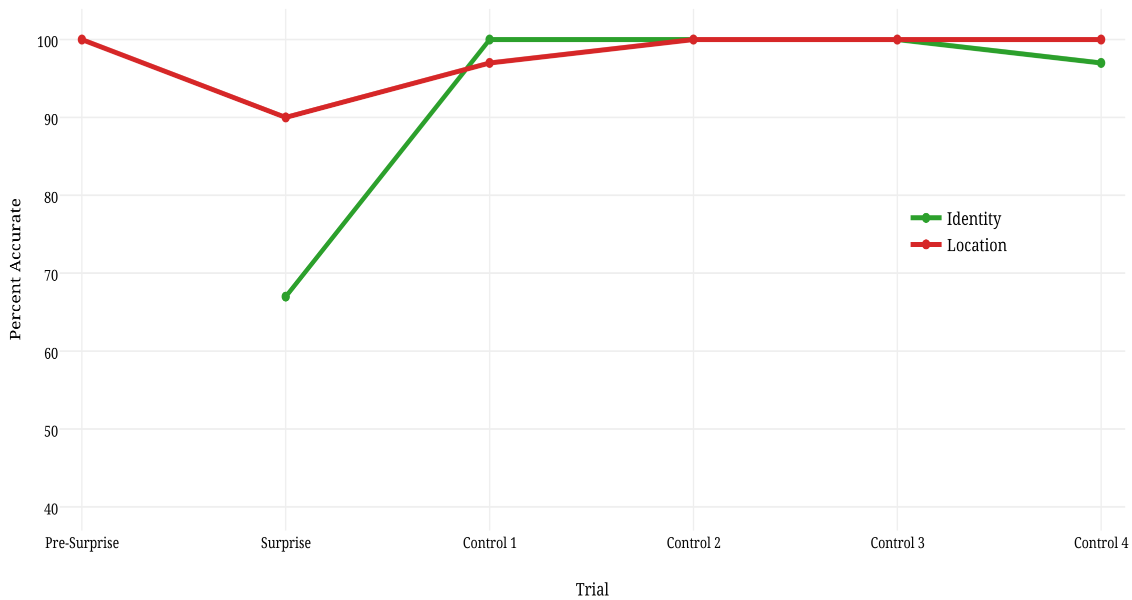

Figure 2.4 Experiment 2 Accuracy Results
Identity and location accuracy compared for each trial. Pre-surprise accuracy is an average from all of the contingency location trials.

Discussion: Experiment 2

In Experiment Two, we were again unable to produce the attribute amnesia effect found by Chen and Wyble (2015). For the attribute amnesia effect, we would have expected to see roughly chance level accuracy (~25%) for letter identification in the surprise trial. This was not the case. We instead saw that participants could correctly identify the letter at a rate of 67%, which is clearly far above chance. This result goes against our hypothesis that the attribute amnesia effect found by Chen & Wyble (2015) was a result of perceiving "letterness" before perceiving the letter. Instead, these findings suggest that interference in the method of reporting may be what is creating the attribute amnesia effect. However, the improvement of accuracy of identity report from the surprise trial to the first control trial suggests an additional interference elicited by asking a surprise question.

General Discussion

3.1 Results

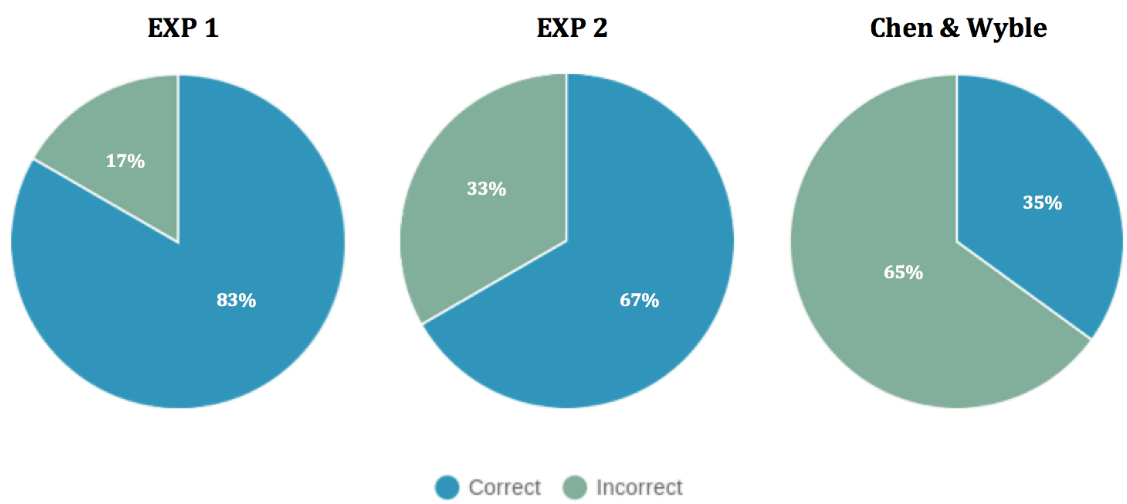

Figure 3.1 Surprise question accuracy of present study and Chen & Wyble (2015) Above chance accuracy of letter identity reporting in the high feature similar stimulus with simplified reporting measures in Experiment One (87%) and the low feature similar stimulus with simplified reporting measures in Experiment Two (67%), compared to the at chance identity reporting of Chen & Wyble's (2015) low feature similar stimulus with un-simplified reporting measures (35%).

Across both experiments we were unable to find any evidence to support attribute amnesia or inattentional blindness. It appears that our participants' ability to correctly identify the target stimulus (the letter) was primarily produced by a reduction of experimental interference in the reporting task. This suggests an important lesson for studies in which perception and memory are unexpectedly probed with retrospective questions. The manner in which the questions are asked might drastically influence the results. Here, when asked to simply click on the location with the mouse and simply type the letter they just saw, participants exhibited well-above chance performance (and thus a lack of inattentional blindness/amnesia). When asked to convert their locations and letter-identity reports to numeric codes before responding, performance dropped to chance (Chen & Wyble, 2015).

Re-examination of Chen & Wyble's Findings

If we now re-examine the results of Chen & Wyble's (2015) findings with the idea of task-based interference in mind, we can start to make more sense of some of the trends that they find.

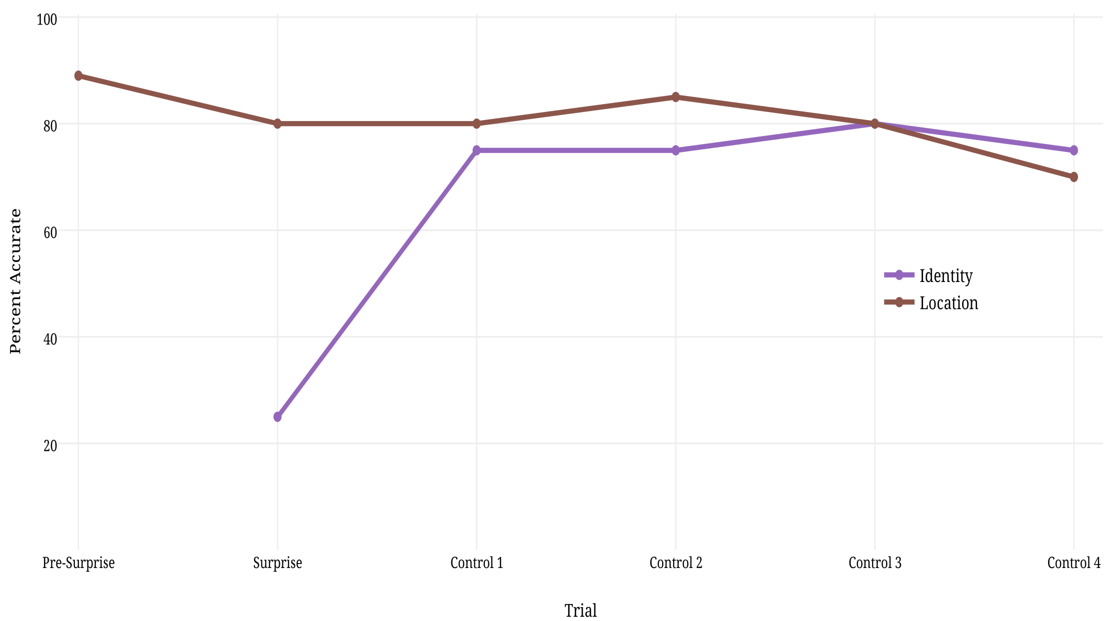

Figure 3.2 Chen & Wyble (2015) Experiment One Accuracy of Report by Trial
Accuracy for Identity and location in each trial of Chen & Wyble's experiment one. The pre-surprise accuracy consists of 155 trials. The surprise trials and control trials are all single trials. Note that the values on the y-axis here are very different than those in Figure 2.3 & 2.4, and that performance on the critical trial (Identity at Surprise trial) is at chance (~25%).

Their first study, in which participants had to undergo the pre-surprise trials 155 times before they are given the surprise question, demonstrated the attribute amnesia effect with a shift of at-chance identity report accuracy on the surprise question to a report accuracy around 80% in the control trials. Chen and Wyble claimed this improvement in support of their finding of the attribute amnesia effect as it suggested that when they were expecting to be questioned on the letter identity, and therefore attended to the identity, the particiapnts were able to accurately report the identity of the letter.

However, based on the findings of our experiments, it would seem that the at chance accuracy in the surprise trial is due to how the question is being asked. What this means is that because of the cognitive demand created by rapid translations in the reporting measures, participants are unable to initially report the identity of the letter, despite having perceived it. A possible explanation for this is that before identity of the letter reaches working memory, the task of responding may be interfering with fragile visual short-term memory. Fragile visual short-term memory is a hypothesized period of four to five seconds of visual representation before the memory either fades or is sustained by attention into working memory (Block, 2011). This fragile stage of memory is believed to be able to contain more visually rich information than working memory, with the estimated capacity of fragile short-term memory being 11 – 12 items, versus working memory capacity being closer to 3 - 4 items. Consequently, some of that information is lost in the transition to working memory (Bronfman et al., 2014). It seems plausible then that Chen & Wyble's method of reporting was inhibiting the amount of information transitioning between fragile visual short-term memory and working memory and thereby resulting in a misleading amnesia report.

While the inhibition of information translation seems to originate from manner of questioning, the sudden shift of task required by a surprise question may compound this effect. They had 155 trials to learn the arbitrary manner of report for location, which had them converting the location to number (e.g., Report "2" if the letter is in the top right). Suddenly, on trial 156, this conversion is interrupted by a new task. This new task requires them to convert the identity of the letter they just saw into a new arbitrary number system (e.g., Report "6" if the letter is "B").

The subsequent leap of report accuracy from the surprise trial to the first control trial can then be explained by participants adjusting to the new task of mental conversions that is now required of them. We see the accuracy rise to the close to 80% that we see for location accuracy for the control trials. What this means is that accuracy in their tasks seems not to be contingent on what part of the stimuli they are attending to, but instead on how trained they are on the task of reporting.

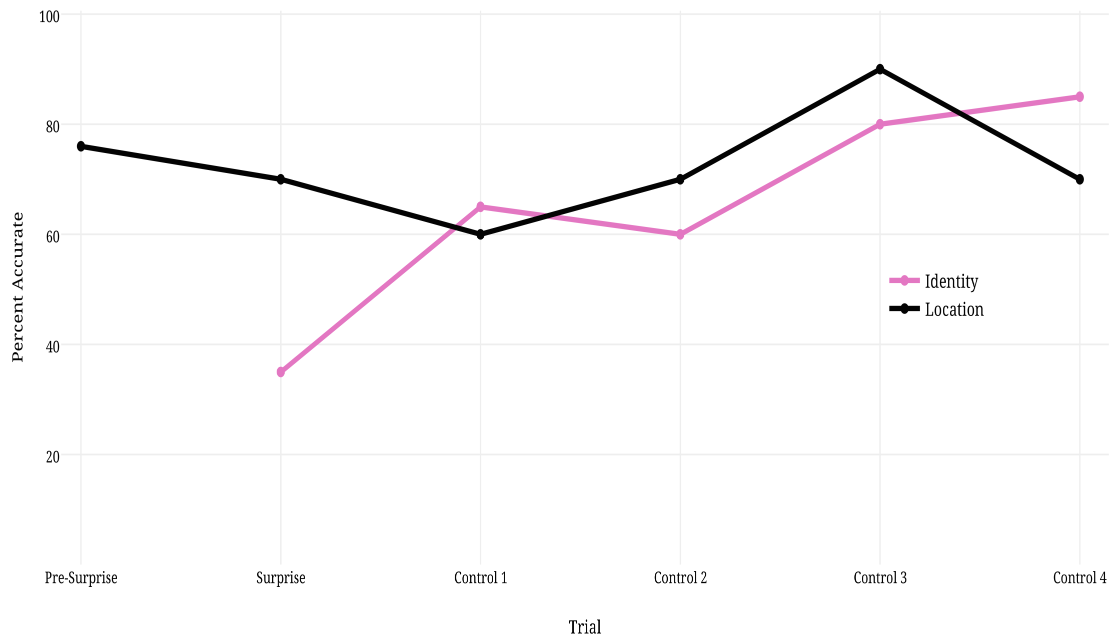

Figure 3.3 Chen & Wyble (2015) Experiment Four Accuracy of Report by Trial Accuracy for Identity and location in each trial of Chen & Wyble's experiment four. In this version of their experiment the pre-surprise question is only asked 11 times.

This notion of training-dependent report accuracy is supported by Chen & Wyble's fourth study. While their Experiment Four findings had similar trends to their experiment one, with roughly 80% accuracy for location report accuracy and near chance (35%) surprise ID report accuracy to an increase to around 70% accuracy across the control trials, there accuracy seemed to be less stable. The location accuracy fluctuated from 60% - 90% in the control trials.

This appears to suggest that by the 11th trial participants haven't been as fully trained in the reporting procedure in comparison to the 155 pre-surprise trials to train from in experiment one, which showed a much more stable accuracy reporting. If this training is the cause of the accuracy fluctuations, it would also suggest that the approximately 80% location accuracy of their Experiment One seems to be the ceiling of reliable accuracy that their reporting measures can achieve.

Cost of Surprise

When we look at our second experiment compared with Chen & Wyble's findings (Figure 3.4) we see similar trends: pre-surprise location accuracy is high, there is a small dip in accuracy for location on the surprises before returning to near pre-surprise accuracy rates during the control trials. This small, but significant, dip may suggest a similar disruption of working memory that resulted in the attribute amnesia effect, but to a lesser degree. In other words, the act of asking a surprise question before asking location again may have been the source of the reduction of report accuracy.

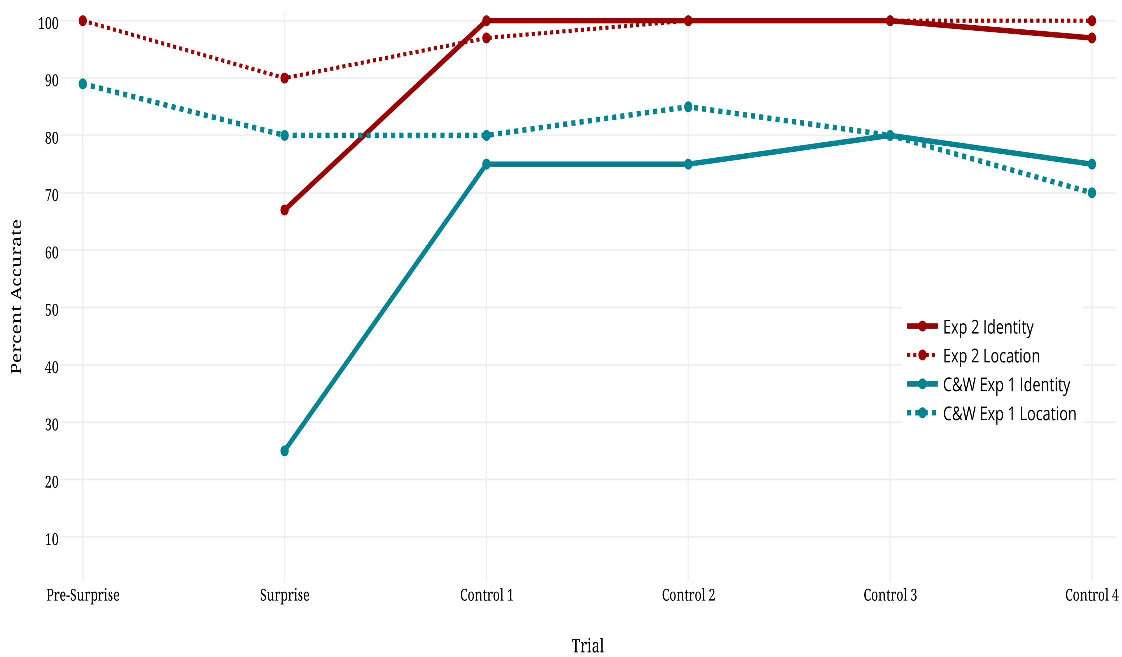

Figure 3.4 Trial Accuracy of Experiment 2 with Chen & Wyble Experiment 1 Comparison of our experiment two (Exp 2) accuracy report for identity and location for each trial to Chen & Wyble's (2015) experiment one (C&W Exp1).

While both location and identity showed a reduction of accuracy because of the surprise question, location reporting seems to be far more resilient to the surprise factor. One likely explanation for this is that it is a trained task, while the surprise question is novel. Additionally, the surprise question involves new instructions that must be read. While the participants were readings this, and before they have fully processed the new question, there may been a loss of letter identity memory associated with figuring out that the new question is asking *for* letter identity. By the time participants figured this out, it was too late and they have already lost the letter identity, while they were still somewhat able to retain the location due to the previous tasks training.

Accuracy Ceiling

As we touched on earlier, another clear difference that arises when we increased the ease of reporting is overall accuracy. Chen & Wyble's Experiment One, despite having 155 trials to train participants in the stimulus and their method of reporting, found only an average of 80% location accuracy across all conditions. On the other hand, in our Experiment Two, participants' accuracy for the location report accuracy for contingent pre-surprise trials was at 100%, with a small drop to 90% accuracy for the surprise trial before returning to near perfect accuracy in the control conditions. Across all trials in our study there was an average of 97% accuracy for location reporting.

A potential implication of this difference of 17% accuracy increase is that the interference caused by Chen & Wyble's numeric-code reporting method is constant enough to create a maximum accuracy less than 100 percent. This would mean that the reporting task is creating an "accuracy ceiling" which may have further implications for their overall results.

3.1 Implications for Future Research

In their present state, the findings of this research are based on one large assumption that has to be cleared up. We are assuming that we can find the attribute amnesia effect if we return to Chen & Wyble's (2015) potentially flawed method of responding. While we feel that this assumption is a relatively safe one to make, as they were able to find their effect in over ten experiments (Chen & Wyble, 2015, 2016), we still must make sure that we can reliably replicate the attribute amnesia effect. This can be easily tested with an exact replication of their version of the experiment.

If we are able to replicate their attribute amnesia findings, the primary implication of this study comes in the form of a call for careful methods of behavioral reporting in memory and perception tasks. Clearly, report accuracy can be easily tampered with by small interferences of cognitive processes. To avoid the likelihood of future false findings it may be prudent to examine the fragility of report-based processes. By doing so, we may create better ways of asking these questions of attention and consciousness, and determine how much weight we should give to inevitable task interference.

"Letterness"

Another potential implication is that of the previously discussed "Letterness". One of the hypotheses we initially explored had to do with perception of categories. The idea was that what might have been creating the attribute amnesia effect was a perception for the category of letters, without having to perceive the individual letters. In other words, if we can perceive the overarching category of a stimulus without identifying sub-categories, we can allocate our attentional resources in a more efficient manner by stopping our perceptual search at the overarching category. For this to be supported, however, we would have expected to find the attribute amnesia effect when we reduced feature similarity in Experiment Two.

This was clearly not the case. However, there was a difference in identity report accuracy on the surprise question between our two experiments (experiment one: 83%; experiment two: 67%). While this difference was not significant, the numeric trends suggest the possibility that the decreased feature similarity of Experiment One forced more attention by inhibiting generic letter perception. This lack of significant findings in this domain may have been due to a simple lack of sample size.

If we look at the differences of accuracy between the surprise trials and control trials of each experiment we do not find a significant change in Experiment One, but we do find one in Experiment Two. As we just addressed, this could be do the overall higher range of accuracies in experiment one in conjuncture with sample size. However, it may be that reduction of categorical distinctions between letters and numbers in Experiment One results participant's memory of the letter being more resilient to the "cost of surprise."

By using our Experiment One as a flange to attach future research to, it should be relatively easy to determine if an ability to perceive “letterness” facilitates the attentional processes of perception.

Careful Questioning

The most important finding of this study is that the processes involved in memory for such attentional tasks is extremely fragile, and as such, we must be incredibly careful about how we structure questions for participants. When we are trying to examine what happens between perception, the memory consolidation of that perception, and the subsequent reporting of that memory we *must* give perception as much benefit of the doubt as possible, as each step towards the ability to report leaves opportunity for undesired interference.

However, as our understanding of the limitations of perception grows we may begin to find a way of looking at phenomena such as change blindness that gives perception that benefit of the doubt. Much of way we see is an ensemble of perceptions, which allows us to digest the world before us at an incredibly efficient rate at the sacrifice of some accuracy. This process happens rapidly and perhaps even without attention, however it is emphasized by attention (Whitney, Haberman, & Sweeny, 2013). We create this ensemble by exploiting redundancy of visual representation to create a largely accurate summary statistic for the world around us (Alvarez, 2011; Cohen, Dennett, & Kanwisher, 2016). This has also been shown to be a process that is constantly adaptive to gradual change as it occurs over time (Albrecht & Scholl, 2010). While summary statistics are an exceptionally powerful and useful tool in day-to-day life, it may result in misleading reports from participants in studies that do not take it into account.

Let’s return to Simons and Levin’s (1997) change blindness study discussed in the introduction. To quickly review, the study had people give directions to another individual who, at some point in the conversation, was obstructed from view and switched out by another individual. They found that many people did not notice this large change. As it is difficult to argue that the participants in this study did not perceive the changed individual, these findings are often attributed to a failure of recollection.

Enter summary statistics.

Similar to inattentional blindness and amnesia studies, change blindness studies have the problematic requirement of the participants being unaware of the target question until it is being asked. As such, if the participant is being asked directions their attention may be on any number of things other than the physical characteristics of the person asking the directions (such as the directions or whatever they were thinking about before the directions). So while they are giving directions, they are creating a summary of what is in front of them (e.g., a person of about average height, male, Caucasian, etc.). When their vision of the person asking for direction is obstructed, and the person subsequently changes, if the summary statistics do not change, they do not notice the person-switch.

In everyday life people don't sporadically change appearance mid conversation. Summary statistics allow us to recreate our past visual memory of the individual to match the impossible incongruences. Because of this, by the time we get around to asking people about what they saw, we are no longer giving perception a chance at being an explanation. As, if someone is reporting that nothing changed, they necessarily must remember the first person as being the person who was swapped in.

Conclusion

The current study initially set out to take steps in understanding inattentional blindness and inattentional amnesia by clarifying findings by Chen & Wyble (2015) which claimed to discover a subset of inattentional amnesia they called “attribute amnesia.”

Our primary hypothesis was a lack of feature similarity between the stimuli used in their paradigm was creating a categorical heuristic that allowed for perception of “letterness” without the perception of letters. We also thought that their method of report collection was unnecessarily complicated and a refinement of it might lead to reporting that was more inline with what the participants were experiencing.

While we were unable to find clear evidence to support the idea of a generic perception of “letterness,” we found that the attribute amnesia effect appeared to be a product of the flawed reporting methods used by Chen & Wyble (2015). It appeared as if by having participants translate letters and locations into arbitrary numbers they were stopping memory consolidation *after* perception had occurred, giving the impression that either perception or an initial memory trace did not occur at all. It is possible that this interruption occurs in a brief window of memory consolidation where the richer perceptual fragile short-term memory is translated into working memory (Block, 2011; Bronfman et al., 2014).

The greater implications of these findings are simple. When researching memory we must make sure that our method of report collection is as straightforward, and natural as possible as to reduce the possible loss of memory of perceived events or objects. The risk of misleading reports, and the ease at which they can be created, are too great to not be as careful as we possibly can be.

References

Albrecht, A. R., & Scholl, B. J. (2010). Perceptually averaging in a continuous visual world: Extracting statistical summary representations over time. *Psychological Science*, *21*(4), 560–567. http://doi.org/10.1177/0956797610363543

Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences*, *15*(3), 122–131. http://doi.org/10.1016/j.tics.2011.01.003

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*(12), 567–575. http://doi.org/10.1016/j.tics.2011.11.001

Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We see more than we can report "cost free" color phenomenality outside focal attention. *Psychological Science*, *25*(7), 1394–1403.

Chen, H., & Wyble, B. (2015). Amnesia for object attributes failure to report attended information that had just reached conscious awareness. *Psychological Science*, 0956797614560648.

Chen, H., & Wyble, B. (2016). Attribute amnesia reflects a lack of memory consolidation for attended information. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(2), 225–234. http://doi.org/10.1037/xhp0000133

Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012). The attentional requirements of consciousness. *Trends in Cognitive Sciences*, *16*(8), 411–417. http://doi.org/10.1016/j.tics.2012.06.013

Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences*, *20*(5), 324–335.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*(1), 1–37.

Halin, N., Marsh, J. E., & Sörqvist, P. (2015). Central load reduces peripheral processing: Evidence from incidental memory of background speech. *Scandinavian Journal of Psychology*, *56*(6), 607–612. http://doi.org/10.1111/sjop.12246

Keizer, A. W., Hommel, B., & Lamme, V. A. F. (2015). Consciousness is not necessary for visual feature binding. *Psychonomic Bulletin & Review*, *22*(2), 453–460. http://doi.org/10.3758/s13423-014-0706-2

King, J.-R., & Dehaene, S. (2014). A model of subjective report and objective discrimination as categorical decisions in a vast representational space. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1641), 20130204–20130204. http://doi.org/10.1098/rstb.2013.0204

Koch, C., & Tsuchiya, N. (2012). Attention and consciousness: related yet different. *Trends in Cognitive Sciences*, *16*(2), 103–105. http://doi.org/10.1016/j.tics.2011.11.012

Lavie, N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, *133*(3), 339–354. http://doi.org/10.1037/0096-3445.133.3.339

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281.

Luck, S. J., Vogel, E. K., & Shapiro, K. L. (1996). Word meanings can be accessed but not reported during the attentional blink. *Nature*, *383*(6601), 616–618. http://doi.org/10.1038/383616a0

Mack, A., Erol, M., Clarke, J., & Bert, J. (2016). No iconic memory without attention. *Consciousness and Cognition*, *40*, 1–8. http://doi.org/10.1016/j.concog.2015.12.006

Mack, A., & Rock, I. (1998). *Inattentional blindness* (1. pbk. ed). Cambridge, Mass.: MIT Press.

Macknik, S., Martinez-Conde, S., & Blakeslee, S. (2010). *Sleights of mind: what the neuroscience of magic reveals about our brains*. London: Profile.

Palmer, J., Boston, B., & Moore, C. M. (2015). Limited capacity for memory tasks with multiple features within a single object. *Attention, Perception, & Psychophysics*, *77*(5), 1488–1499. http://doi.org/10.3758/s13414-015-0909-2

Persuh, M., Gomez, M., Bauer, L., & Melara, R. D. (2014). Feature-based inattentional blindness: loss of awareness to featural information in fully attended objects. *Attention, Perception, & Psychophysics*, *76*(8), 2229–2239. http://doi.org/10.3758/s13414-014-0712-5

Pöppel, E., Held, R., & Frost, D. (1973). Residual visual function after brain wounds involving the central visual pathways in man. *Nature*, *243*(5405), 295–296. http://doi.org/10.1038/243295a0

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

Ruz, M., Worden, M. S., Tudela, P., & McCandliss, B. D. (2005). Inattentional amnesia to words in a high attentional load task. *Journal of Cognitive Neuroscience*, *17*(5), 768–776.

Shapiro, K. L., & Luck, S. J. (1999). The attentional blind: A front-end mechanism for fleeting memories. *Fleeting Memories*, *17*, 95–118.

Simons, D. J. (1996). In sight, out of mind: When object representations fail. *Psychological Science*, *7*(5), 01–305.

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

Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, *1*(7), 261–267. http://doi.org/10.1016/S1364-6613(97)01080-2

Sünram-Lea, S. I., Foster, J. K., Durlach, P., & Perez, C. (2002). Investigation into the significance of task difficulty and divided allocation of resources on the glucose memory facilitation effect. *Psychopharmacology*, *160*(4), 387–397. http://doi.org/10.1007/s00213-001-0987-9

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136. http://doi.org/10.1016/0010-0285(80)90005-5

Vandenbroucke, A. R. E., Fahrenfort, J. J., Sligte, I. G., & Lamme, V. A. F. (2014). Seeing without knowing: Neural signatures of perceptual inference in the absence of report. *Journal of Cognitive Neuroscience*, *26*(5), 955–969. http://doi.org/10.1162/jocn_a_00530

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

Whitney, D., Haberman, J., & Sweeny, T. D. (2013). From textures to crowds: multiple levels of summary statistical perception. *The New Visual Neurosciences*. *MIT Press;*

Wolfe, J. M. (1999). Inattentional amnesia. *Fleeting Memories*, *17*, 5.

Wolfe, J. M., & Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research*, *37*(1), 25–43. http://doi.org/10.1016/S0042-6989(96)00111-3