2015

Exploring the Long-Term Consequences of False and Correct Recognition

Matthew T . Sabia
Lehigh University

Exploring the Long-Tem Consequences of False and Correct Recognition

by

Matthew T. Sabia

A Thesis

Presented to the Graduate and Research Committee

of Lehigh University

in Candidacy for the Degree of

Master of Sciences

in

The Department of Psychology

Lehigh University

May 1, 2015

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Human memory is not a precise picture of the past and thus is prone to error. It is susceptible to the presentation of distorting information (e.g. Loftus, 1977; Loftus, Miller, & Burns, 1978), and people sometimes struggle to distinguish between highly similar alternatives (Guerin et al., 2012). However, certain situations seem to foster excellent memory even when the distinctions to be made are quite difficult (Brady et al., 2008). I conducted three experiments in an attempt to better understand how avoidance of common memory errors influences long-term memory. All three experiments asked participants to study individual objects during encoding. Immediately after, they performed a visual recognition test, with some conditions fostering false recognition (Guerin et al., 2012). Forty-eight hours later, a second recognition test was administered to test the long-term effects of initial correct and false recognitions. Experiment 1 asked whether rejecting a non-target foil of high similarity to a previously encountered target at Test 1 might lead to richer encoding of that foil, leading to better recognition of that foil at Test 2. Experiment 2 asked whether rejecting the foil initially might cause subjects to mistake the representation of the foil for the target representation. I found evidence against these hypotheses. Correct rejection of high-similarity foils at Test 1 led to subsequent failure to recognize those foils at Test 2 (Experiment 1). Moreover, even when the targets were again presented at Test 2 (Experiment 2), initial correct rejections most often led to subsequent failure to recognize targets. Experiment 3 asked whether correct rejections at Test 2 were made on the basis of impaired recognition for conceptual or perceptual details. The results suggest that both conceptual and perceptual details play a role in Test 2 rejections. These findings are discussed with respect to misleading post-

event information, reconsolidation, intentional forgetting, and the role of sleep in memory consolidation. Overall, the results of these studies produce a coherent pattern. Rejections on an initial test often lead to subsequent memory failure, while initial false recognitions of a target-related foil often lead to later false recognition of the same foil, but do not necessarily interfere with access to the target representation.

Introduction

Human memory does not provide an exact reproduction of the past. It is a reconstruction of the past formed from pieces of information from many different sources. While we are usually able to piece together the past in a useful manner through this reconstruction process, memories also highly susceptible to distortion (Bartlett, 1932). For example, consider the story of Ronald Cotton, who was wrongfully convicted of the rape of Jennifer Thompson based on the victim's eyewitness testimony, testimony that she insisted upon in spite of the presence of her actual assailant at a second trial (Hansen, 2001). Several different types of memory distortions and errors have been studied, such as the creation of new false memories, specifically creation of false memories after imagining events (e.g. Mazzoni & Memon, 2003) or objects (e.g. Lyle & Johnson, 2006) or by automatic activation of words that are highly associated with target words (e.g. Roediger & McDermott, 1995). Additionally, it has been shown that schemas influence retrieval processes (e.g. Bartlett, 1932; Bransford & Franks, 1971; Anderson & Pichert, 1978) and that additional information can be integrated into memories upon their reactivation (e.g. Hupbach, Gomez, Hardt, & Nadel, 2007). Most important for the present work are studies of gist-based false recognition, a robust memory phenomenon in which people often mistakenly recognize items that are similar but not identical to previously encountered items; that is, false recognitions occur because the previously encountered item and the newly presented one share gist information (e.g. Koutstaal & Schacter, 1998; Guerin, Robbins, Gilmore, & Schacter, 2012). Also of relevance are studies demonstrating that misleading post-event information distorts memory for details of an event, as these studies involve the possibility of introducing plausible but false information into an episodic memory (e.g. Loftus, 1977; Loftus, Miller, & Burns, 1978; for a review see Loftus, 2005).

These latter two types of memory distortions share some important similarities, but also have one key difference. In examinations of the misinformation effect, subjects view an event and are then presented with misleading information in a narrative following the event. However, they are led to believe that the information in the narrative is true, thus creating a divergence between what subjects saw and what they were subsequently informed was true. Meanwhile, in paradigms designed to investigate gist-based false recognition, subjects study a series of target items and are then tested on both targets and high-similarity foils (which I will refer to as target-related foils), but are not led in any way to believe that foil items are in fact targets; on the contrary, they are warned about the potential confusion and make errors in spite of it. This makes the implications that stem from this paradigm even more powerful, because these false recognitions are committed in spite of subjects' awareness of their susceptibility to false recognitions.

In what follows, I will review the relevant research on these two memory processes, and present pilot data that extend a paradigm for investigating gist-based false recognition developed by Guerin et al. (2012) as background to the research questions that my Master's thesis addresses. The main question that my thesis addresses is how both false recognition and the avoidance of false recognition modulate subsequent long-term memory. To answer these questions, I investigated the long-term consequences of accepting vs. rejecting target-related foils in a gist-based false recognition paradigm. I also examined the possibility that a target-related foil could be encoded more richly than other non-targets, and that the process of reactivation that takes place or the fluency of processing of a target-related foil may even cause the original target memory to be overwritten. I also conducted an exploratory study on the role of conceptual and perceptual details in the observed poor target memory that was displayed, in order to better understand the causes underlying rejections at Test 2. Taken together, these studies sought to advance the understanding of memory error and memory modification by outlining a set of conditions under which both initial memory failures and initial successful memory can lead to subsequent memory errors.

I will begin by describing the misinformation effect, a memory phenomenon that broadly aligns with the topic of this thesis. I will then review memory processes that may underlie this effect (namely, source monitoring difficulties and memory reconsolidation). I will then relate misinformation-based memory error to related work in visual recognition memory. Of particular importance in this section is the discussion of the paradigm used by Guerin et al. (2012) mentioned above, as my thesis attempts to replicate but also to critically extend this paradigm by adding long-term memory tests. Finally, I will present the results of pilot work that form the basis of the research questions I address in this thesis.

The Misinformation Effect

Loftus (1977) conducted an early study examining the influence of misleading post-event information on memory for details of an event. In her study subjects viewed slides detailing an accident in which a red car strikes a pedestrian at a crosswalk. In one scene, a green car drives past the accident scene without stopping. Subjects then answered a series of questions about the slides, with half of the subjects receiving the critical question, “Did the blue car that drove past the accident have a ski rack on the roof?” This question implies that the car that drove past the accident was blue rather than green. The other half of subjects received the same question with the word “blue” omitted. After a filler task, subjects were then given a color wheel and asked to point out the color that best matched specific objects from the scene. Loftus found that only 4 out of 48 misled subjects selected a color on the wheel that closely matched the correct color of the passing car, as opposed to 13 of 47 subjects who were not misled. Additionally, misled subjects chose colors in the blue spectrum significantly more often than subjects who were not misled. As a follow-up, Loftus then had the control subjects return for another test a week later. These subjects completed the same questionnaire they had completed initially, except that half of them now received the misleading critical question. Loftus found that subjects who were misled on the second test often shifted their initial selection closer to a shade of blue, but overall their selections did not correspond to a blue shade as often as the one in subjects who received misleading information in the first test. Loftus interpreted this result as showing increased difficulty in shifting people’s memory after having committed to a selection (see also Experiment 3 in Loftus, 1977). Loftus, Miller, and Burns (1978) extended this research to show that misleading information not only leads to subtle changes in color recognition, but to the acceptance of conceptual detail changes and other perceptual elements as well. Using the same set of slides, the critical slide was now one where the red car stopped at a traffic sign. For half of the subjects, the red car stopped at a stop sign; for the other half, the car stopped at a yield sign. In Experiment 1, subjects completed a questionnaire about the slides immediately after viewing them; for half of subjects, one question indicated that the car had stopped at a stop sign, and for the other half, the same question indicated that the car had stopped at a yield sign. This was organized in such a way that half of subjects received consistent information from the questionnaire while the other half received misleading information. When presented with a two alternative forced-choice recognition test, misled subjects chose the correct scene only 41% of the time, in comparison to 75% correct choices in the consistent group, indicating that the misleading information that contained a conceptual change, interfered with subjects' memory.

A follow-up study eliminated demand characteristics as an explanation. It could be argued that subjects believed that they were "supposed" to respond in the recognition test using information from the questionnaire. In this experiment, subjects completed the same experimental design as the one discussed above, except one-third of the subjects received a version of the critical questionnaire item that did not reference a stop sign or a yield sign. Also, upon completion of the experiment, subjects were asked to indicate what sign they remembered seeing and what they remembered being referenced on the questionnaire. Subjects who had been misled performed significantly worse on the critical item in the recognition memory test than subjects who received consistent or irrelevant information regarding the sign. Additionally, subjects who chose incorrectly on the recognition test were significantly worse at correctly identifying the critical information from both the slide and the questionnaire when asked after the experiment. Only 12% of subjects who responded incorrectly on the recognition test reported differences in what they believed they had seen and what the questionnaire said, indicating that demand characteristics play a minimal role. Another follow-up found that presenting misleading information just prior to test increased the impact of the misinformation compared to when misleading information was presented immediately after viewing the event, with a delay before testing. Loftus et al. interpret this effect in terms of time. They argue that when misinformation is presented immediately after viewing an event, its impact is strongest after presentation and decays temporally in a manner consistent with the decay of the event itself. As a result, misinformation effects appear when a test occurs immediately after misinformation presentation, but its impact is lessened when the test is delayed. However, when misinformation is presented after a delay, the trace for the original event has faded, making the impact of the misinformation more powerful.

The idea of memory trace decay over time, especially with regard to misleading information, will become more relevant in the discussion of my pilot studies, as the presence or absence of a pre-test delay proves to be a critical element in the primary phenomenon of interest in this thesis. Trace decay has largely been studied in working memory as it relates to the ability to retain information over short spans of time (e.g. Baddeley & Scott, 1971) and has often been pitted against interference theories with regard to explaining forgetting. While interference theories posit that competing memories may block retrieval of information and render it temporarily inaccessible (e.g. Skaggs, 1933), trace decay theories suggest that a consolidated memory trace may become weaker over time and may be completely forgotten (Thorndike, 1913). Recent views of trace decay have pointed to the adaptive value of losing information from memory (e.g. Rosenzweig, Barnes, & McNaughton, 2002) and the failure of interference theories to explain the inability to retain rote information in memory (Cowan et al., 2004). Additionally, Hardt, Nader, and Nadel (2013) have posited that trace decay is the primary cause of the loss of hippocampus-dependent (i.e., consolidated) memories because of the efficient pattern separation processes that take place in that region of the brain. Meanwhile, they propose that interference is a better explanation of memory failures that take place during active processes such as explicit encoding and retrieval (i.e., memories that are either not fully consolidated or have been reactivated; see section on reconsolidation) and is also more likely when information is stored in brain areas with poor pattern separation.

What Causes The Misinformation Effect?

Loftus and Palmer (1974) interpreted the misinformation effect as reflecting changes in the original memory, an interpretation that eventually came across stiff opposition. McCloskey and Zaragoza (1985) argued that the findings of Loftus' experiments did not show change in the original memory, as the paradigm used in her experiments was unfit to detect such changes. They reasoned that because the misleading information was presented as an option in the subsequent memory test, subjects who did not remember the original information would be biased to choose the misleading item. To test whether the original memory was altered by the post-event misleading information, they developed a modified test procedure in which the subsequent test contained the item from the original slides (e.g., green car) and a novel item that crucially differed from the misleading information (e.g., a yellow car). McCloskey and Zaragoza contended that if the original memory was actually impaired by the misleading question, subjects would not be able to distinguish between the original item and a novel one. In a series of experiments, they found no differences between the misled and control groups in their memory for the details of the event, supporting their hypothesis that memory for the original event was not impaired. The issue of what could be measured using different testing procedures and the appropriateness of each would become an issue of contention among memory researchers in the coming years.

In an attempt to reconcile the existing literature on the effects of misleading information on memory, Tversky and Tuchin (1989) introduced yet another testing procedure, a yes/no recognition test which required people to respond to possible alternatives individually (original items, misleading items, and new items). Misled and control subjects were equally good at rejecting new foils, but misled subjects were poorer at recognizing the original information than control subjects. They took this as evidence that misleading post-event information did in fact influence the original representation in memory.

To what degree the presentation of misleading information impairs the original memory is a question that remains unresolved. Is it simply a matter of not being able to distinguish the source of information? Or is the original memory trace actually changed by the presentation of the misleading information? Is some combination of these alternatives possible? Several accounts on the nature of memory change offer their own explanations for the misinformation effect. Most important for the purpose of this thesis are the source-monitoring framework and the reconsolidation account of memory change. The explanations that these accounts offer can also be applied to visual recognition memory and the ability to make difficult memory distinctions.

Source Monitoring

The source-monitoring framework posits that people make rapid judgments about the source of activated memories based on qualitative characteristics such as the amount of perceptual detail available (Johnson, Hashtroudi, & Lindsay, 1993). These judgments can be influenced by a variety of factors that lead to error. If the initial memory was encoded well, and if the person is not under any time pressure at the time of the source judgment, these judgments are often highly accurate. However, if the initial trace was encoded poorly, or there are stressors present at the time of the source decision (e.g. cognitive load or time pressure), the source judgment is often less accurate. Importantly, source judgments can be made based on heuristic or systematic judgments. Heuristic judgments are often rapid, general criteria for determining the source of information (e.g., whether a certain piece of information is familiar or not) and are most commonly used for source judgments. Systematic judgments are less frequent, slower judgments made based on more specific criteria (e.g., how different is this item from an item that has definitely been encountered before?).

A number of studies have examined the role of source monitoring in the misinformation paradigm. For example, Lindsay and Johnson (1989) found that subjects who were given a source memory test after viewing a series of slides and reading a narrative containing misleading information were much less likely to report having actually seen misleading information than those who simply took a recognition test; that is, they correctly attributed the source of the misleading information to the narrative. However, another study by Lindsay (1990) indicated that misleading information could still impact people's memory when source distinctions are difficult. Subjects viewed slides and were presented with misleading information about them, either immediately after the slides in the same context as the learning environment ("difficult" condition), or immediately before the test (which was 48 hours later) in a different context from the learning environment ("easy" condition). At test, subjects were given an explicit warning that the information given in the post-event narrative was wrong, but only subjects in the "easy" condition seemed able to avoid using information from the narrative when asked about the original slides. These findings suggest that misinformation is much less influential when distinctions between sources of information are easy. Additionally, this finding regarding the timing of misinformation presentation stands in contrast to Loftus et al. (1978) who found that misinformation presented after a delay had the greatest impact. This may indicate the importance of the role of context. In Lindsay's study, misinformation presented prior to the delay also was presented in the same location, while misinformation presented after the delay was always presented in a different location from the learning context, a factor not manipulated by Loftus et al. Taken together, these studies show that misinformation may have a greater impact before a delay under certain conditions and after a delay under others. Additionally, the timing of misinformation presentation can interact with other factors such as contextual match to influence memory.

Reconsolidation

Research on memory reconsolidation has shown that memories can be altered after encoding, supporting the original hypothesis by Loftus (1975). The reconsolidation hypothesis states that every time a memory is reactivated, it re-enters a fragile state in which it can be modified. A study by Nader, Schafe, & Le Doux (2000) reintroduced previously explored ideas of reconsolidation with regard to fear memories. They found that fear memories in rats could be erased by reactivating the original fear memory after consolidation, and then injecting a protein synthesis inhibitor. Work by Hupbach, Gomez, Hardt, and Nadel (2007) extended these ideas about memory malleability to human episodic memory. In their study, subjects learned a list of objects on Day 1, and then were presented with a new set of items forty-eight hours later. Prior to the presentation of new items, subjects were either asked to recall the instructions they received on Day 1 or not. An additional 48 hours later, all subjects were asked to recall the list of items learned on Day 1. Subjects in the reminder group mistakenly reported List 2 objects as part of List 1 ("intrusions") significantly more often than the no reminder group, suggesting that the reminder successfully reactivated the consolidated memory trace for List 1 items. This suggests that learning a second list while the first list was activated altered memory for the original list. The findings show that presenting new, task-relevant information can result in memory change under certain conditions.

Hardt, Einarsson, and Nader (2010) have suggested that memory reconsolidation underlies the misinformation effect. They point to the similarities between the experimental paradigms used to examine both reconsolidation and the misinformation effect as arguments for the viability of the reconsolidation hypothesis in explaining the misinformation effect. In particular, they argue that the process at the heart of the misinformation effect is retrieval-induced plasticity, the idea that a memory becomes malleable upon reactivation (as seen in Hupbach et al., 2007).

The above hypothesis regarding the role of reconsolidation in the misinformation effect received support by Chan, Thomas, and Bulevich (2009). In their primary study, subjects viewed a video portraying a fictional terrorist attack and then were either given a cued recall task or completed a filler task. After a delay, all subjects listened to an audio narrative that contained some misleading details about the original video and were subsequently given a final memory test. They found that subjects who did not complete the first recall test and therefore did not reactivate the original memory did not suffer memory impairment for misleading details in the audio narrative; however, subjects who did complete the first recall test were highly susceptible to misinformation based error. These results corroborate Hardt et al.'s (2010) idea that reconsolidation processes may underlie the misinformation effect.

In summary, the misinformation effect literature clearly shows that misleading information changes memory in some way, regardless of whether the distortions reflect retrieval (source memory errors) or storage processes (actual memory impairment). However, the misinformation and reconsolidation paradigms are not the only ways to examine how memory for learned information might be influenced by the introduction of new information. Another example is the paradigm developed by Guerin et al. (2012) discussed earlier. They developed this paradigm in order to bring together research reporting faulty recognition memory (e.g. the misinformation effect) as well as work that has found a massive storage capacity for visual detail in human recognition memory (see below). In spite of the differences between the misinformation effect paradigm and the Guerin et al. paradigm (which have already been discussed), the false recognition paradigm provides a useful alternative for testing the impact of potentially misleading information on memory. The discussion below further illuminates the origins of this paradigm and illustrates the findings of a key study that forms the foundation of the studies discussed in this thesis.

Gist-based False Recognition

When misinformation is presented before a test of episodic memory, people often report falsely remembering the misleading information at test. This effect has long been used to show that episodic memory is prone to error. However, research also exists showing that human capacity to remember specific details is quite large (Shepard, 1967; Standing, 1973). In both Shepard's and Standing's work, however, subjects were not asked to make difficult distinctions when their memory was tested; it should have been quite easy for subjects to perform well, even if they only retained gist information about studied items and not specific details.

Brady, Konkle, Alvarez, and Oliva (2008) were interested in examining the storage capacity of human visual memory and in particular the ability to hold enough detail to make difficult distinctions between old and new information. In their study, subjects viewed 2500 pictures of real world objects for three seconds each. Immediately after the study phase, their memory was tested by a two-alternative forced-choice test in which a studied object and an object that was either completely novel, an exemplar of the same category, or the same object in a different state were presented. The results showed that subjects correctly recognized targets on 92% of trials in the novel condition, 88% of trials in the exemplar condition, and 87% of trials in the state condition. Brady et al. interpreted these results in support of an extremely large capacity of human memory to store details, as evidenced by extremely high accuracy rates under test conditions that required retrieval of highly specific detail. Although previous work had suggested that humans generally have a large memory capacity, no prior experiment had directly pitted images that were learned in large volume against high-similarity alternatives that required detailed recall for accurate selection. These reports of a massive detail capacity in memory also stand in contrast to longstanding reports of poor recognition memory from paradigms such as the misinformation effect. Importantly, these findings also stand in contrast to one of the key arguments from McCloskey and Zaragoza (1985) against the idea of memory change in the misinformation effect paradigm. They argued that when subjects simply forget the details of the original event, they are led to make a biased guess (in the "original" testing format utilized by Loftus and colleagues). However, the above research suggests that the specific details are not usually forgotten. Thus, the primary source of bias that McCloskey and Zaragoza claim is introduced by the "original" testing format does not appear to be a concern. This may be a reason to suspect that response bias does not play a large role in misinformation effect findings, which would support Loftus' hypothesis of memory alteration.

In an attempt to reconcile the reports of a large capacity for detail in human memory such as those made by Shepard (1967), Standing (1973), and Brady et al. (2008) with reports of episodic memory failure, Guerin et al. (2012) developed an experimental paradigm designed to understand the conditions under which recognition memory might be flawed and under which conditions it might be accurate. In their study, subjects studied a series of 144 items for 500 ms each, and were then presented with a three alternative recognition test. Figure 1 outlines the conditions of this test.

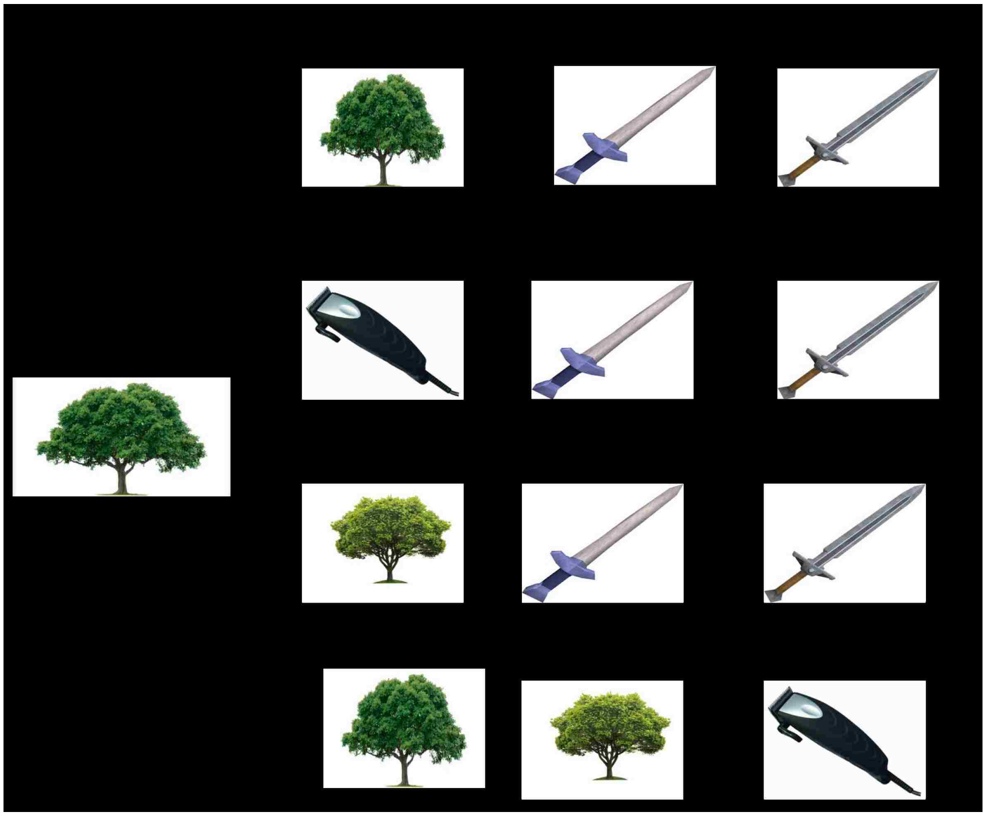

Figure 1. Illustration of encoding and test conditions in Guerin et al.'s (2012) study. Items were counterbalanced such that different images were used in all conditions.

Two of the three items were always exemplars of the same category ("tokens" of the same "type"; see Kanwisher, 1987 for use of this terminology in a similar context) and shared a common verbal label. Subjects could choose one of the three items as a target, or they could reject all items. In the *baseline target* condition, a target item was presented along with two conceptually unrelated (different "type") foils. In the *baseline foil* condition, no target was present, nor were any "tokens" of a target "type". In the *single related item* condition, a "token" of a target "type" (target-related foil) was presented alongside two different "type" foils. Finally, in the *target and related item* condition, a target and a target-related foil were presented alongside one different "type" foil (see Figure 1 for a graphical example of the possible choices in each condition). The results showed that subjects were quite accurate in identifying the target when unrelated distracters were present (73%) and were able to correctly reject when no alternatives shared an item type with a target (69%). Both of these numbers are lower than ones reported by Brady et al. (2008) under similar testing conditions; however, this discrepancy can likely be attributed to the presence of additional options in Guerin et al.'s (2012) design. However, subjects frequently committed gist-based false recognition (that is, they selected the item that was a different "token" of the same "type" as the target) in the single related item condition (41%), but gist-based false recognition dropped in the target and related item condition (13%). These results suggest that in the presence of the target, gist-based false recognition is significantly more unlikely when the target is present than when it is absent.

One question that remained unanswered following this study is the mechanism behind this phenomenon. Is excellent detail recognition due to target reinstatement itself or the increased attention to details? A follow-up study added a two related items condition in which two target-related foils were presented. If people are actually making responses based on attention to detail, they should not commit gist-based false recognition in the two related items condition. People should examine both related foils and determine that neither is a target. Conversely, a high rate of gist-based false recognition in this condition would indicate that target presence is the primary factor in allowing people to overcome their susceptibility to false recognition. Subjects were still susceptible to gist-based false recognition in the two related items condition (46%), indicating that it is the presence of the target itself that is responsible for the ability to make difficult distinctions, rather than increased attention to detail.

When directly comparing the Guerin et al. (2012) paradigm to the misinformation effect paradigm, the acceptance of the target-related foil in false recognition appears to be akin to the acceptance of a misleading detail in the misinformation paradigm. However, no studies have questioned the impact of a false recognition on later memory. Indeed, the central claim of the misinformation effect is that misleading information leads to a memory change as manifested on a subsequent test. Therefore, it is imperative to investigate the long-term consequences of gist-based false recognition in order to understand the full impacts of memory errors.

Long-Term Consequences of False Recognition?

For my thesis, I am interested in how conditions that foster false recognition affect the long-term retention of target items. In particular, I was interested in whether accepting or rejecting a target-related foil during an initial memory test affects the target memory formed during learning, especially after a delay, as a way of linking the literature on visual recognition memory to long-term memory phenomena such as the misinformation effect. Returning to the Guerin et al. (2012) paradigm, subjects are "misled" by the presentation of the target-related foil, and accepting that foil as a target is quite similar to the integration of misleading information into an episodic memory. However, acceptance of misleading information tends to lead to long-term memory impairment; as I will discuss shortly, my pilot study found that this was not the case in using the Guerin et al. paradigm. I assume that utilizing the false recognition paradigm allows for a more nuanced analysis of the memory processes that sometimes lead to long-term acceptance of misleading information, because the paradigm established by Guerin et al. (2012) eliminates demand characteristics as a possible explanation for results. Subjects in this paradigm are warned about high-similarity foils at test, whereas errors in misinformation effect studies are based on subjects' belief that the misinformation presented to them is in fact true. Additionally, a second recognition test allows for a detailed analysis of the nature of long-term memory impairment. By investigating how subjects respond at a second recognition test based on their responding at the first test, I can directly test the impact of memory errors on subsequent memory and provide a framework for making predictions about specific ways in which memory error might modulate long-term memory.

Pilot Studies

In Pilot Study 1, I asked how false recognition impacts subsequent long-term memory by extending Guerin et al.'s paradigm with a second, delayed recognition test. This second test was designed with McCloskey and Zaragoza (1985)'s "modified" procedure for testing the impact of post-event misinformation in mind. To understand whether false recognition indeed led to memory impairment, original target items had to be pitted against alternatives that had not been seen before. According to McCloskey and Zaragoza, this provides an unbiased measure of whether or not long-term memory had been impaired by misinformation. Subjects in this pilot study first encoded a series of images and were tested under identical conditions to the Guerin et al. study. Then, subjects returned to the lab 48 hours later for a second recognition memory test. In this test, subjects were presented with a target from the encoding phase and a new "token" of the target "type" that had not been encountered before. As in the first test, subjects were asked to select the image that had been displayed during encoding and had the option to respond “all new” if neither image had been shown during encoding.

I predicted that, after false recognition in the *single related item* condition at Test 1, subjects would have poorer memory for associated targets at Test 2 than after correct rejection in the *single related item* condition. This prediction was based on the reconsolidation account discussed earlier in the Introduction. I assumed that presentation of the target-related foil reactivated the representation of the associated target. I also assumed that subsequent acceptance of the target-related foil would then lead to impairment of the representation of the associated target in memory. By contrast, correct rejection of the target-related foil should not impair memory for the associated target. Different predictions could be made based on the source monitoring account. This account would predict that, because there is no potential source confusion at Test 2, subjects should identify targets regardless of their responses at Test 1.

Proportions of Test 2 responses to targets used in the *single related item* condition at Test 1 are displayed in Figure 2. Following false recognition in the *single related item* condition, subjects were able to recognize corresponding targets at Test 2 more often than following correct rejections. Additionally, “all new” responses at Test 2 were more prevalent following correct rejections than following false recognition. This pattern of results was contrary to the above predictions and presented an interesting puzzle. How could correctly rejecting a foil item at an initial test impair long-term memory for the corresponding target item? In order to understand this finding further, it was necessary to examine the basis on which correct rejections in the *single related item* condition at Test 1 are made. It is possible that correct rejections at Test 1 are based on good target memory; as a result, the target-related foil is rejected because it differs from the target representation. It is also possible that correct rejections are the result of poor target memory; in this case, Test 1 rejections are made indiscriminately, but on these particular trials produce correct responses. In order to examine these possibilities, a second pilot study was conducted in which the 48-hour delay between Test 1 and Test 2 was removed. The differences in Test 2 selection based on Test 1 responses in the *single related item* condition observed in the initial pilot were not present in this second pilot (see Figure 3). This indicated that good target memory is the basis for correct rejections in the *single related item* condition at Test 1, as access to target representations remained intact when the delay between tests was removed.

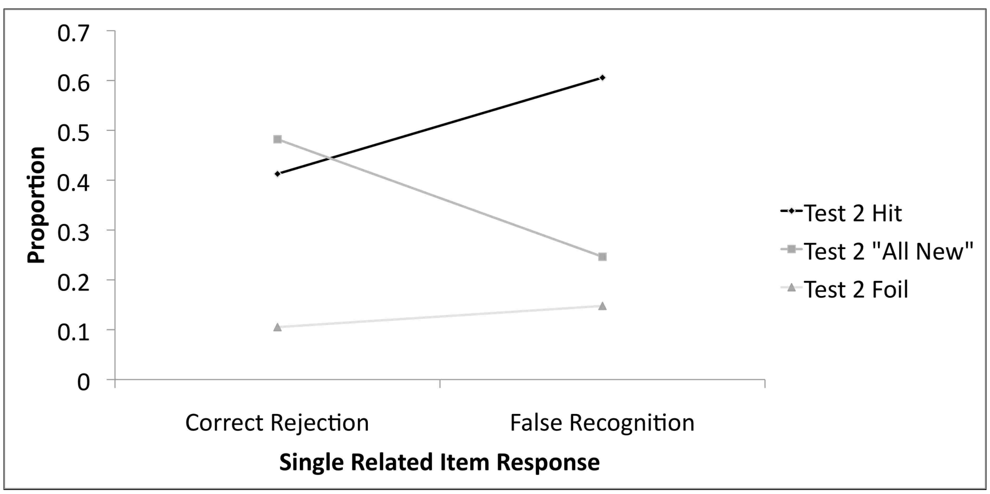

Figure 2. Test 2 response proportions based on Test 1 responses on *single related item* trials in Pilot Study 1.

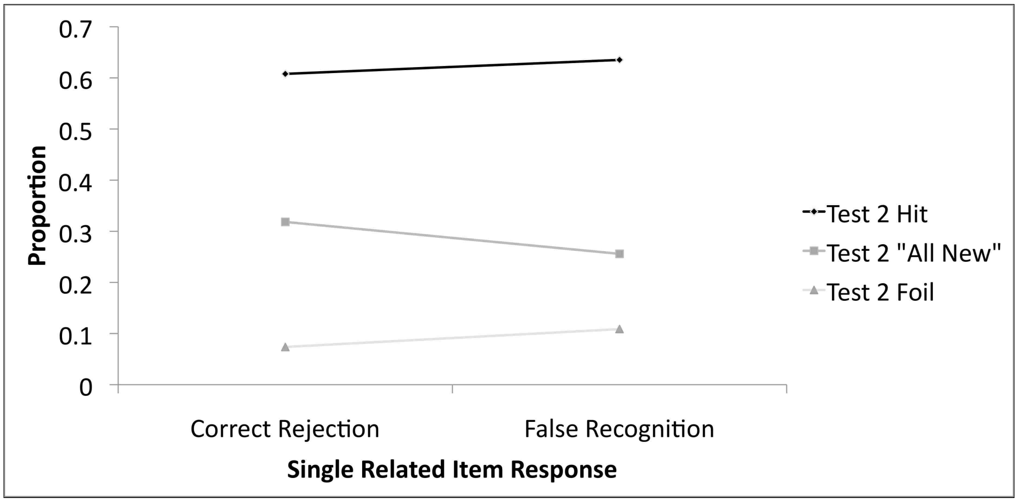

Figure 3. Test 2 response proportions based on Test 1 responses on *single related item* trials in Pilot Study 2 (no 48-hour delay between tests).

In the following section, I describe three experiments that test several explanations for why people might reject targets following correct rejections of target-related foils. Experiment 1 tested whether correct rejections of the target-related foil at Test 1 may lead to richer encoding of the foil due to the difficulty in rejecting a foil that shares a “type” with a target item (see Craik & Lockhart, 1972 for more on the impact of levels of processing effects on long-term memory). Experiment 2 tested whether the target-related foil is mistaken for the target as a result of the foil being processed very fluently (see Whittlesea & Williams, 2001), or because correct rejections may reactivate the target representation such that the foil representation may be mistaken for the target representation. This secondary hypothesis aligns with the reconsolidation account discussed in the Introduction. As established in my second pilot, correct rejections are the result of identifying the discrepancy between a target and a target-related foil. Therefore, it is possible that a correct rejection necessitates a reactivation of the target representation that is then replaced by the presented foil representation. Finally, Experiment 3 tested the level to which conceptual and perceptual detail influenced Test 2 rejections. Episodic memory for perceptual and conceptual detail can differ on a number of important dimensions such as rate of decay (Brainerd & Reyna, 2002), which was discussed briefly in the Introduction. Taken together, these studies sought to investigate how both false recognition and the avoidance of false recognition on an initial test might modulate long-term memory.

Experiment 1: Enhanced encoding of the related foil?

Experiment 1 tested whether or not correct rejections in the *single related item* condition at Test 1 led to enhanced encoding of the target-related foil, thus enabling it to interfere with attempts to retrieve the target representation. This hypothesis proposes that the deeply encoded and therefore better-remembered foil (Craik & Lockhart, 1972) competes for retrieval. Rejecting the foil potentially requires richer encoding of its perceptual details than accepting it based on gist information. Additionally, the proposed stronger foil encoding for correct rejections might cause target memory impairments after a delay because of different forgetting rates for the target and the target-related foil. Memory for the target might decline more rapidly than memory for the foil, which was viewed more recently and was perhaps encoded more richly. This would mean that the perceptual details of the target might be further along in the decay process than the details of the target-related foil, making those details more accessible.

In order to test the hypothesis that the target-related foil is particularly well encoded following an initial correct rejection of the foil, the format of Test 2 was altered from my pilot studies. In the *single related item* and *target and related item* conditions, subjects were presented with the target-related foil seen during Test 1 alongside a new "token" of the same "type", while in the *baseline target* and *baseline foil* conditions subjects were presented with the originally encoded target alongside a new "token" of the same "type" (for the purposes of the design of this study, I will refer to previously encountered items collectively as "old tokens", and new foils as "new tokens"). Subjects were asked to select any items that they had seen in any prior phase of the experiment (encoding or Test 1). If target rejection in Test 2 of my pilot study is based on enhanced memory for the target-related foil, then subjects should recognize the target-related foil more often following correct rejections than false recognition in the *single related item* condition.

Methods

Subjects and Design

Thirty-two undergraduates (20 females, 12 males, mean age = 18.50 years, *SD* = 0.80) were recruited from the Lehigh University participant pool and participated in the study for course credit. The experimental design was a 4 x 2 within subjects design. Factor 1 was Test 1 condition with 4 levels: *baseline target*, *baseline foil*, *single related item*, *target and related item*. Factor 2 was Test 2 condition with 2 levels: "old tokens" and "new tokens".

Materials

The stimulus material consisted of 1196 color photographs representing everyday objects that could be readily named. Always three images represented different versions of the same object (e.g., three different kinds of trees, three different kinds of swords, etc.). The majority of these images were the same stimuli used by Guerin et al. (2012). Some of Guerin et al.'s images were replaced with images found on the Web (images.google.com). During encoding, 144 distinct object images were presented. At Test 1, sets of three images were presented at a time, with two "tokens" representing the same "type" (e.g., oak tree and elm tree) and one image of a different "type" (e.g., electric razor). This test had four conditions that are outlined in Figure 1. Seventy-two of the 144 target images were assigned to the *single related item* condition, in order to maximize the number of trials in this condition. Twenty-four of the target images were assigned to the remaining three conditions. In the *baseline target* condition, a target image was presented alongside two new images that shared a "type", but were different "tokens". In the *baseline foil* condition, target images were omitted, and only new images were presented. In the *single related item* condition, a target-related image (same "type", different "token") was presented alongside a pair of new images shared a "type" but were different "tokens". Finally, in the *target and related item* condition, a target image and a target-related image were presented alongside a new image. At Test 2, pairs of images were presented. For targets assigned to the *baseline target* and *baseline foil* conditions, a target was presented alongside a new foil of the same "type". For targets assigned to the *single related item* and *target and related item* conditions, the target-related foil shown at Test 1 was presented alongside a new foil of the same "type". Additionally, Test 2 also contained 24 new pairs of items. There were a total of 168 Test 2 trials.

Procedure

Encoding. The encoding procedure was a replication of the procedure used by Guerin et al. (2012). Subjects viewed 144 color photographs one at a time for 500 ms each. For each image, subjects were asked to indicate whether the item presented would fit into a 13-inch box in the real world, with a box of that size provided for reference. Subjects were explicitly instructed to remember the items for a later memory test.

Test 1. After completing a short (~5 minute) filler task, Test 1 was administered. Again replicating Guerin et al.'s procedure, three images were always presented at a time (for conditions, see Materials). The subject's task was to select the image that exactly matched an image from the study phase. If none of the images could be recognized as a previously studied item, subjects were instructed to reject all items. Responses in this test were not time-constrained.

Test 2. Subjects returned to the lab for a second recognition test 48 hours later. In this test, subjects were always presented with pairs of images. For the 144 experimental trials, one of these images was either a target image from the encoding phase or a target-related foil encountered during Test 1. Targets were shown when they had been assigned to the *baseline target* or *baseline foil* conditions, and target-related foils were shown when the associated target had been assigned to the *single related item* or *target and related item* condition. The second image was always a new "token" of the same "type" that had not been encountered before. For the 24 control trials, pairs of images were shown that shared a "type", but no "token" of this "type" had been displayed as a target or displayed during Test 1. Subjects were asked to select the item they had encountered before in the experiment (either in the encoding phase or during Test 1). As in Test 1, if none of the images could be recognized, subjects were instructed to reject both items. Responses were not time-constrained.

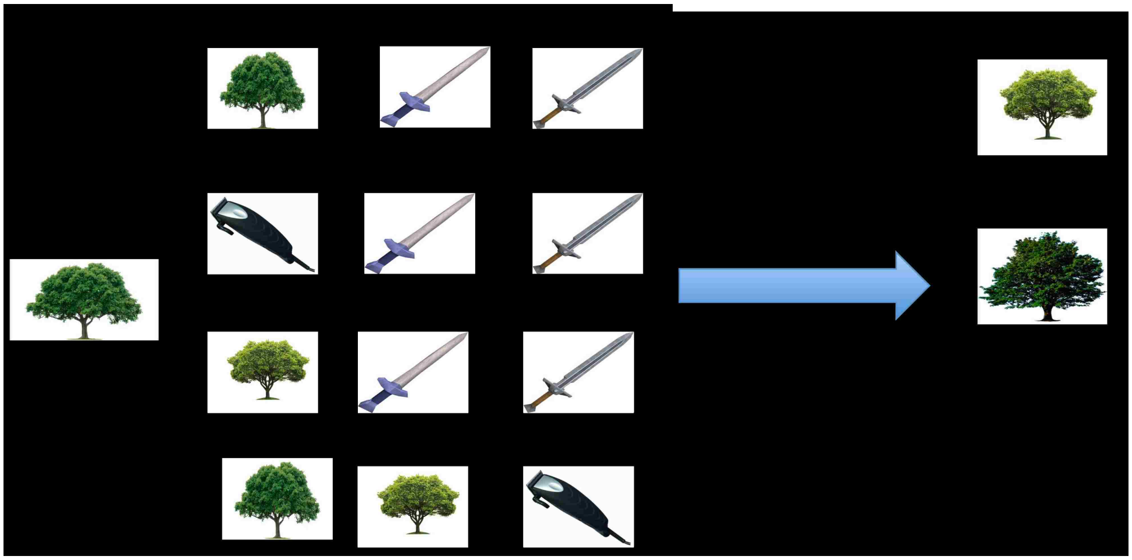

Figure 4. Illustration of conditions in Experiment 1. Items were counterbalanced such that different images were used in all conditions at Test 1.

Results

Test 1

Test 1 results are depicted in Figure 4. Following Guerin et al., I first examined how often subjects committed false alarms to conceptually independent lures in the *single related item* (gist-based false recognition) and *baseline foil* conditions. This analysis provides an indication of whether gist-based false recognition is a unique phenomenon compared to false alarms to standalone foils. Subjects chose the target-related foil in the *single related item* condition significantly more often than they chose the standalone item in the *baseline foil* condition, $t(31) = 11.21$, $p < .001$, indicating that the target-related foil was chosen due to its shared gist with a target item; a standard gist-based false recognition effect. Next, I compared gist-based false recognition rates in the *single related item* and *target and related item* conditions. This analysis provides an indication of whether gist-based false recognition is more common when the target is absent. False recognition was significantly more prevalent in the *single related item* condition than in the *target and related item* condition, $t(31) = 10.97$, $p < .001$, indicating that gist-based false recognition was more prevalent when the target was absent. I then compared, within *the target and related item* condition, gist-based false recognitions versus false alarms to the standalone foil as an indication of whether the conceptual relationship between the target-related foil and the target made it a particularly attractive foil. False recognition rates in the *target and related item* condition were indeed significantly higher than false alarm rates to the standalone foil in this condition, $t(31) = 3.30$, $p = .002$, indicating that the shared gist between the target and target-related foil is responsible for selections of the target-related foil. Additionally, *baseline target* accuracy was significantly higher than accuracy in the *target and related item* condition, $t(31) = 2.62$, $p = .014$.

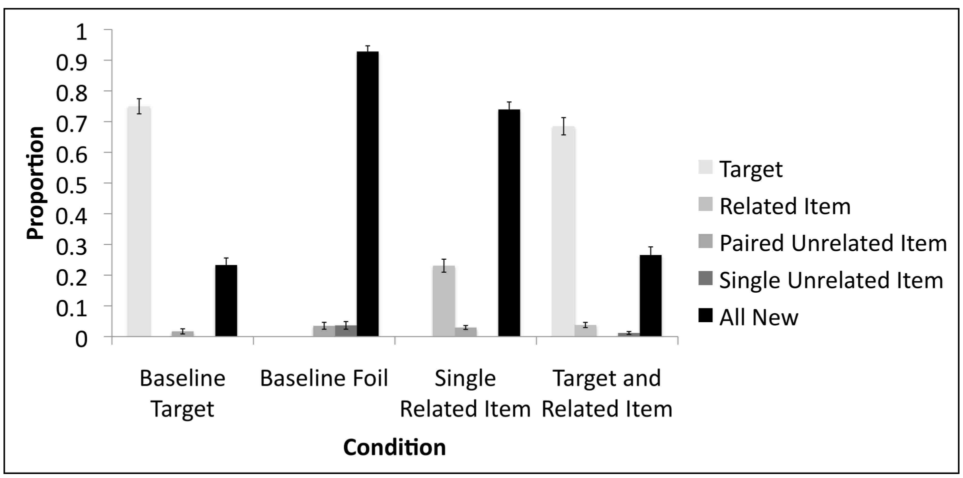

Figure 5. Response proportions at Test 1 across conditions in Experiment 1.

Test 2

I performed three analyses on Test 2 data in order to examine my questions of interest. I first ran two separate t-tests comparing the corrected number of recognized images (hits minus false alarms) between two pairs of conditions: the *baseline target* and *baseline foil* conditions (in which targets were displayed at Test 2), and the *single related item* and *target and related item* conditions (in which target-related foils were displayed at Test 2). This was done in order to isolate differences in selection of "old tokens" versus "new tokens" at Test 2 based on the experimental conditions. I then conducted a one-way repeated measures ANOVA comparing frequencies of "all new" responding across all five Test 2 conditions to know whether misses differed across conditions, and to know whether correct rejections are particularly high when new pairs of images are presented at Test 2. Here, the analysis was not split because "all new" was an option in all Test 2 conditions. Finally, focusing solely on trials corresponding to the *single related item* condition at Test 1, I conducted a 2 x 2 repeated measures ANOVA examining the frequency of Test 2 hits and misses based on whether subjects correctly rejected or falsely recognized target-related foils at Test 1. This will allow for direct comparison of how Test 1 responses influence Test 2 responses in the *single related item* condition while still analyzing the data within the ANOVA framework. I predict that subjects will choose the target-related foil at Test 2 more often following correct rejections than false recognitions. All analyses were conducted on raw data, but proportions are depicted in graphs for ease of interpreting the data.

First I will focus on the t-tests for corrected recognition. A paired-samples t-test comparing numbers of correctly identified targets in the *baseline target* vs. the *baseline foil* conditions was significant, $t(31) = 15.89$, $p < .001$, indicating that targets were more often correctly recognized in the *baseline target* than the *baseline foil* condition. This seems to indicate that the repeated presentation and selection of the target in the *baseline target* condition boost Test 2 recognition of the target. A second paired-samples t-test comparing the numbers of correctly identified target-related foils in the *single related item* vs. *target and related item* conditions was also significant, $t(31) = 5.30$, $p < .001$, indicating that more foils were recognized in the *single related item* condition than the *target and related item* condition.

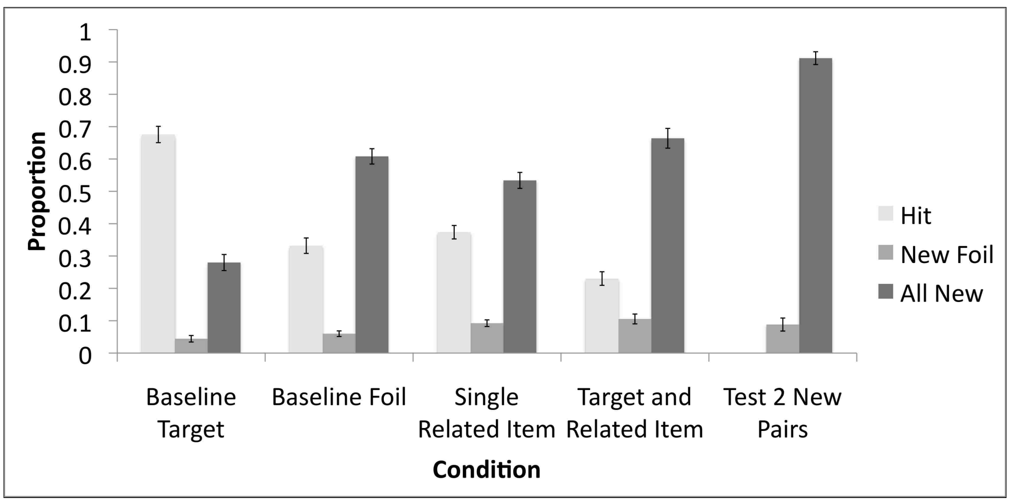

Figure 6. Proportion of Test 2 responses across conditions in Experiment 1. “Hit” in the *baseline target* and *baseline foil* conditions corresponds to correctly identified targets. “Hit” in the *single related item* and *target and related item* conditions corresponds to correctly identified foils from Test 1.

For the one-way ANOVA comparing “all new” responses across all five Test 2 conditions, Mauchly’s test of sphericity was violated, so Greenhouse-Geisser corrected values are reported. There was a significant main effect of Test 1 condition on “all new” responses, F (2.88 89.37) = 129.23, $MSE = 10.67$, $p < .001$, $\eta^2 = .807$. Post-hoc comparisons using Bonferroni-corrected p values showed significant differences between all conditions ($ps < .05$), except the *target and related item* and *baseline foil* conditions. These findings reflect (a) that correct rejections were high in the *all new* condition, and (b) that "all new" responses are a mirror image of correct identifications. This also aligns with findings from my pilot work indicating that the new foil at Test 2 was an extremely rare response.

Crucially, I further analyzed Test 2 responses following *single related item* trials at Test 1. Specifically, I looked at the differences between hits and "all new" responses for items that were correctly rejected vs. falsely recognized at Test 1 (see Table 1 and Figure 7). I analyzed Test 2 responses following false recognition and correct rejections with a 2 (Test 1 response: correct rejection vs. false recognition) x 2 (Test 2 response: hit vs. "all new" responses) repeated measures ANOVA. False alarms at Test 2 (that is, responses to the new foil) were not included in this analysis because they were extremely rare (means presented in Table 1). This analysis revealed a main effect of Test 1 response, $F\ (1, 31) = 103.83$, $MSE = 84.23$, $p < .001$, $\eta^2 = .770$, indicating that there were overall more correct rejections than false recognition in the *single related item* condition at Test 1. There was also a main effect of Test 2 response, $F\ (1, 31) = 11.12$, $MSE = 81.19$, $p = .002$, $\eta^2 = .264$, indicating that there were more "all new" responses than hits overall. Importantly, there was an interaction between Test 1 and Test 2 responses, $F\ (1, 31) = 117.75$, $MSE = 42.68$, $p < .001$, $\eta^2 = .792$. Simple effect analyses revealed that there were significantly more "all new" responses than hits at Test 2 for images that had been correctly rejected at Test 1, $F\ (1, 31) = 46.43$, $p < .001$, $\eta^2 = .600$, and significantly more hits than "all new" responses for images that were falsely recognized at Test 1, $F\ (1, 31) = 58.91$, $p < .001$, $\eta^2 = .655$.

	False Recognitions	Correct Rejections	T1 False Alarms
Hits	11.19	15.19	0.53
"All New"	3.97	33.03	1.44
T2 False Alarms	1.47	5.03	0.16

Table 1. Mean number of target-related foils recognized (hits) vs. "all new" responses and false alarm rates based on responses in the *single related item* condition at Test 1 in Experiment 1. Data from the false recognition and correct rejection columns will be displayed in Figure 7.

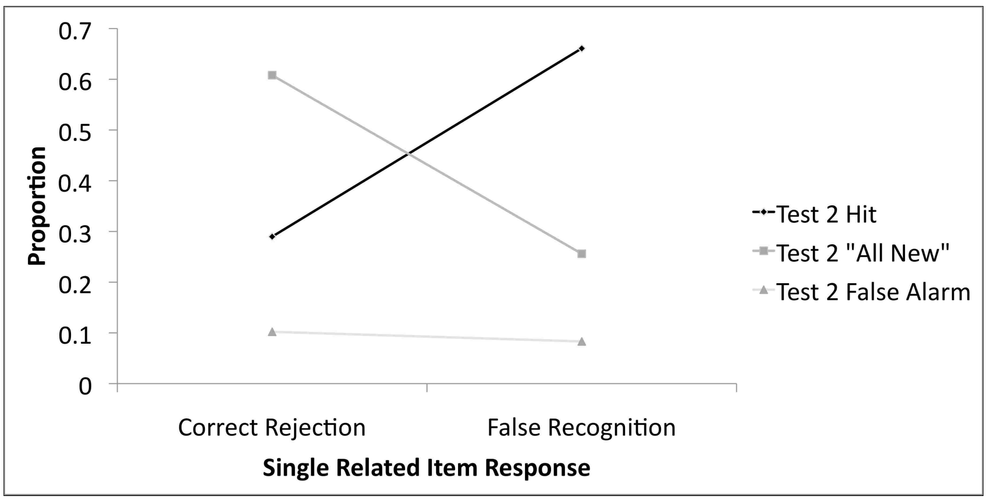

Figure 7. Proportion of Test 2 responding based on responses in *single related item* condition at Test 1 in Experiment 1. "Hits" in this graph represents responses to target-related foils.

Discussion

Looking first at Test 1 results, the pattern of effects established by Guerin et al. (2012) were replicated. The target-related foil was more often selected in the *single related item* condition than the standalone foil in the *baseline foil* condition, which constitutes a significant gist-based false recognition effect. False recognitions in the *single related item* condition occurred significantly more often than false recognitions in the *target and related item* condition, indicating that the presence of the target as an option reduced false recognition rates. However, false alarms to target-related foils in the *target and related item* condition occurred significantly more often than false alarms to unrelated foils in the same condition, indicating that while false recognition effects were reduced in the *target and related item* condition, they were not completely eradicated.

Looking at corrected recognition in Test 2 based on Test 1 conditions, corrected recognition of targets in the *baseline target* condition was better than in the *baseline foil* condition. This was an unsurprising finding given that the target image had been displayed twice in the *baseline target* condition and only once in the *baseline foil* condition. This aligns with work on instance theories in memory (e.g. Logan, 2002) that posit that memory for an item is strengthened simply by encountering the item multiple times. Additionally, target-related foils were more often selected in the *single related item* condition than in the *target and related item* condition. This seems to indicate that the likelihood of selecting the target-related foil at Test 1 modulated memory for those foils at Test 2. Selecting the target-related foil was more common in the *single related item* condition than in the *target and related item* condition at both Test 1 and Test 2, regardless of the change in task demands at Test 2.

For the *single related item* condition, I predicted that correct rejections in comparison to false recognition would result in richer encoding of and therefore better memory for the target-related foils. However, the results provide strong evidence against this hypothesis. Correct rejections in the *single related item* condition at Test 1 resulted in fewer hits and more “all new” responses, while the opposite was true following false recognition. Ultimately, a pattern of results emerges that resembles misinformation effect studies using the original test (e.g. Loftus, Miller, & Burns, 1978). Subjects that accept the target-related foil were more likely to later recognize that image than those who were able to reject it. However, in order to get a complete picture of how this task aligns with tasks investigating the impact of misleading post-event information, the original and misleading information must be pitted against each other (see Experiment 2).

An alternative explanation for the results in Experiment 1 could be that rejection of the foil during Test 1 leads to a tendency to repeat the same response on the second test. Evidence from work with educational materials on multiple-choice tests has found that selection of a foil at an initial test often leads to production of that foil on a later test (Roediger & Marsh, 2005). Additionally, the simple act of encountering a particular test lure increases subjective ratings of truth with respect to that lure (Toppino & Brochin, 1989). Returning to the present work, it seems possible that the act of correct rejection in the *single related item* condition could become an episodic memory trace in and of itself, such that when the target-related foil is presented again at Test 2, it cues the response (“reject”) given at the previous test. Likewise, the act of having chosen a particular item could be encoded into memory and then cued by presentation of that same item.

Taken together, the pattern of results from this experiment indicate that both the frequency with which an image is encountered and the response given to an image play a role in Test 2 recognition. This aligns with findings that the attention that is paid to an item modulates long-term memory for it, particularly with regard to top-down attention (Uncapher & Rugg, 2009; Uncapher, Hutchinson, & Wagner, 2011). This makes sense in the context of false recognitions, which would seem to engage more top-down processes of attention (since committing a false recognition in the *single related item* condition, by definition, represents a failure to sufficiently investigate the perceptual details of the image).

Thus far, I have established that correct rejection of target-related foils impairs subsequent memory for the target. This relative impairment is eliminated when the second memory test takes place immediately after the first test, suggesting that subjects utilize their memory for the initial targets to reject highly similar foils. Experiment 1 indicates that memory for the rejected foil is not enhanced by the act of rejecting it. Rather, it appears that selecting the foil at Test 1 increases the likelihood of selecting it again at Test 2, aligning with past work on the impact of the presentation of misleading post-event information (Loftus, 1977; Loftus, Miller, & Burns, 1978). In contrast, correctly rejecting target-related foils at Test 1 leads to failure to recognize both items that were rejected and targets that are associated with them. This brings into question the basis of Test 2 rejections. Understanding why subjects reject previously encountered images at Test 2 may clarify the exact impact of initial correct rejections on long-term memory modulation. Based on these results, it is possible that subjects are rejecting based on the perceptual details of presented images (that is, rejecting specific "tokens"), or that subjects are rejecting entire concepts (that is, entire "types"). Experiment 3 will provide a more detailed discussion of this question.

Experiment 2: Foil Representation Mistaken for the Target?

Experiment 2 tested whether correct rejections of the target-related foil cause subjects to later assume that the foil was the target, or in other words, whether they "replace" the target with the foil in memory over time. One possible reason for this phenomenon might be the overall effect of processing fluency on memory. Target-related foils were encoded more recently, which may result in more fluent processing of the foil item, which in turn leads subjects to believe that this item is in fact a target (see Whittlesea & Williams, 2001, for a discussion of the impact of processing fluency on memory). Alternatively, based on the reconsolidation account discussed earlier one might propose another possibility. One could argue that perhaps it is only following correct rejections in the *single related item* condition that the original target is reactivated sufficiently to allow for replacement by another alternative. Given that Pilot Study 2 established that correct rejections at Test 1 in the *single related item* condition are the result of target retrieval, it is possible that this retrieval makes the target representation susceptible to modification by the target-related foil.

To test whether target-related foils were sometimes mistaken for targets, targets and target-related foils were presented together in Test 2 of Experiment 2, and subjects were asked to indicate which item was part of the *original* memory set. Thus, in contrast to Experiment 1, Test 2 pitted targets and target-related foils against each other, with the instruction to select targets only. This provided a clear examination of what is represented in memory when two corresponding "tokens" of a particular "type" were pitted against each other. I expected subjects to be able to respond accurately on the second test if the target and foil had already been presented together on the first test (*target and related item* condition), because subjects should quite often have been able to correctly identify the target at Test 1. If the target and target-related foil had not been presented together previously (*single related item* condition), and the foil takes the place of the target over the period of 48 hours, then subjects should incorrectly choose the foil as the target.

Methods

Subjects and Design

Thirty-two subjects (18 females, 14 males, mean age = 20.12 years, *SD* = 2.05) received $15 for their participation in this study. The experimental design was a 4 x 2 within subject design. Factor 1 was Test 1 condition with 4 levels: *baseline target*, *baseline foil*, *single related item*, *target and related item*. Factor 2 was Test 2 condition with 3 levels: target vs. target-related item, target vs. new item, and target and target-related item absent.

Materials

The materials were identical to those used in the pilot studies and Experiment 1.

Procedure

Encoding. The encoding phase was identical to that of the pilot studies and Experiment 1.

Test 1. Test 1 was identical to the pilot studies and Experiment 1.

Test 2. As in the previous study, subjects returned to the lab 48 hours after Test 1 for a second recognition test, and again they were presented with pairs of images. For the 144 experimental trials, one of the images viewed by subjects was always a target from the encoding phase. For targets corresponding to the *baseline target* condition at Test 1, the second image was an image of a different “type” relative to the target. This was done to alleviate any potential confusion with regard to the instructions, as displaying different “tokens” of the same “type” on every trial may lead to confusion. For targets corresponding to the *baseline foil* condition at Test 1, the second image was a different "token" of the same "type" relative to the target image. For targets corresponding to the *single related item* condition and *target and related item* conditions at Test 1, the second image was always the target-related foil that had been displayed during Test 1. For the 24 control trials, the two images displayed were of different "types", and neither "type" had been displayed at encoding or at Test 1. Subjects were instructed to select the item they had seen during the encoding phase. As in the pilot studies and Experiment 1, if subjects could not recognize either image as being displayed during encoding, they were instructed to reject both images, and responses on each trial were not time-constrained.

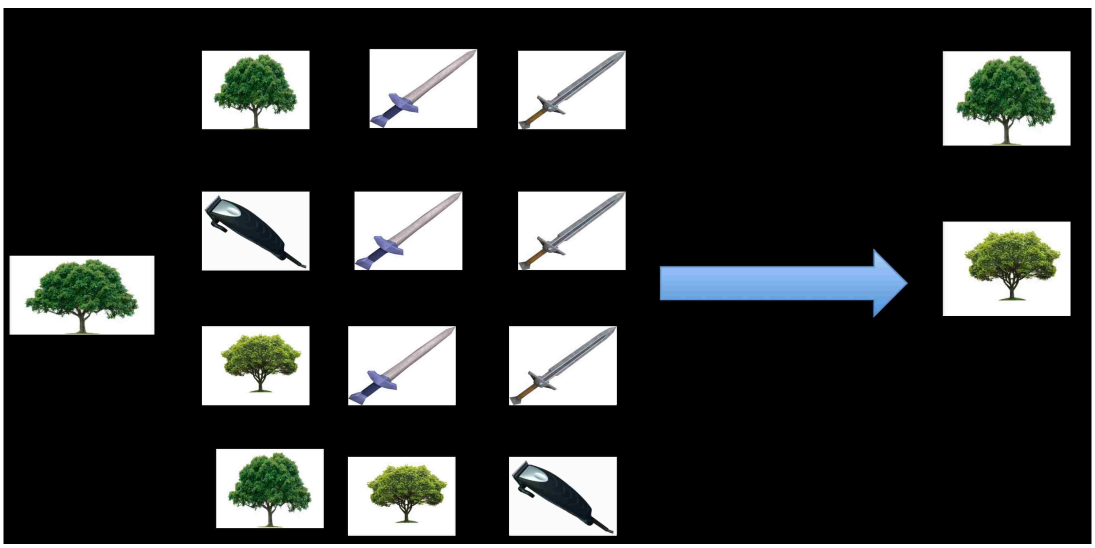

Figure 8. Illustration of conditions in Experiment 2. Items were counterbalanced such that different images were used in all conditions at Test 1.

Results

Test 1

Test 1 results are depicted in Figure 7. Subjects chose the target-related foil in this condition significantly more often than they chose the standalone item in the *baseline foil* condition, $t(31) = 11.15$, $p < .001$, indicating a significant gist-based false recognition effect. They also committed false recognitions significantly more often in the *single related item* condition than in the *target and related item* condition, $t(31) = 11.29$, $p < .001$. However, false recognition rates in the *target and related item* condition were significantly higher than false alarm rates to the standalone foil, $t(31) = 3.59$, $p = .001$, and *baseline target* accuracy was again significantly higher than accuracy in the *target and related item* condition, $t(31) = 2.20$, $p = .035$, indicating that gist-based false recognition effects were reduced but not eliminated in the *target and related item* condition.

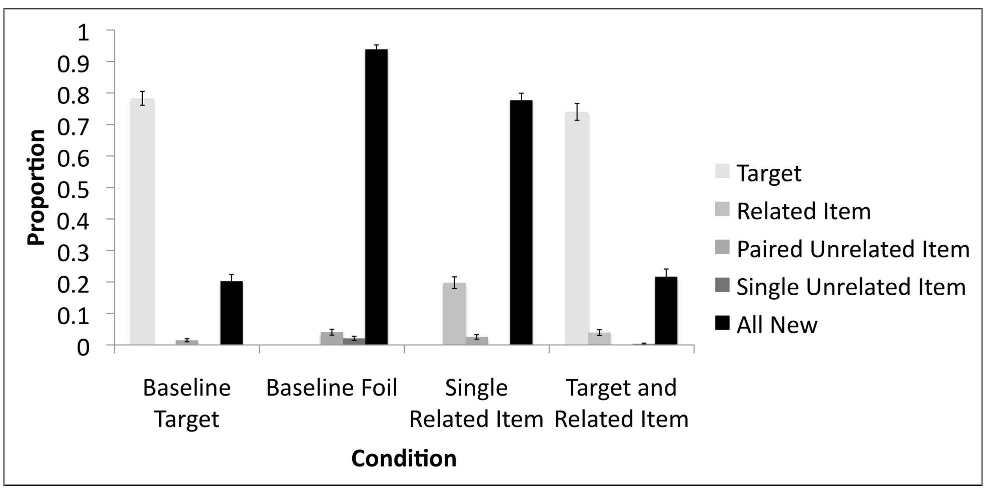

Figure 9. Response proportions at Test 1 in Experiment 2 across conditions.

Test 2

I performed three analyses on Test 2 data in order to examine my questions of interest. I first conducted a one-way repeated measures ANOVA comparing the corrected number of recognized images across all four Test 2 conditions that contained target images. This was done in order to test how well subjects recognized targets at Test 2 based on the condition the target had been assigned to in Test 1, thus informing whether the presence or absence of the target at Test 1 impacted Test 2 responding. I then conducted a one-way repeated measures ANOVA comparing frequencies of “all new” responding across all five Test 2 conditions to know whether misses differed across conditions, and to know whether correct rejections are particularly high when new pairs of images are presented at Test 2. Finally, focusing solely on trials corresponding to the *single related item* condition at Test 1, I ran a 2 x 3 repeated measures ANOVA examining the frequency of Test 2 hits and misses based on whether subjects correctly rejected or falsely recognized target-related foils at Test 1. This will allow for direct comparison of how Test 1 responses influence Test 2 responses in the *single related item* condition while still analyzing the data within the ANOVA framework. I predict that subjects will select targets more often following Test 1 false recognitions compared to Test 1 correct rejections. All analyses were conducted on raw data, but proportions are depicted in graphs for ease of interpreting the data.

For the one-way ANOVA examining corrected recognition, Mauchly’s test of sphericity was violated, so Greenhouse-Geisser corrected values are reported. There was a significant main effect of corrected recognition across Test 1 conditions, F (1.47, 45.48) = 21.67, MSE = 76.02, $p < .001$, $\eta^2 = .411$. Post-hoc comparisons using Bonferroni-corrected p values revealed significant differences between all conditions (ps < .05) except between the *baseline target* and *target and related item* conditions, and between the *single related item* and *baseline foil* conditions. Indeed, the pairs of conditions that did not differ can be defined by whether or not target images had been presented during

Test 1 in those conditions, indicating that target presence at Test 1 led to good target recognition at Test 2, while target absence led to significantly poorer Test 2 recognition.

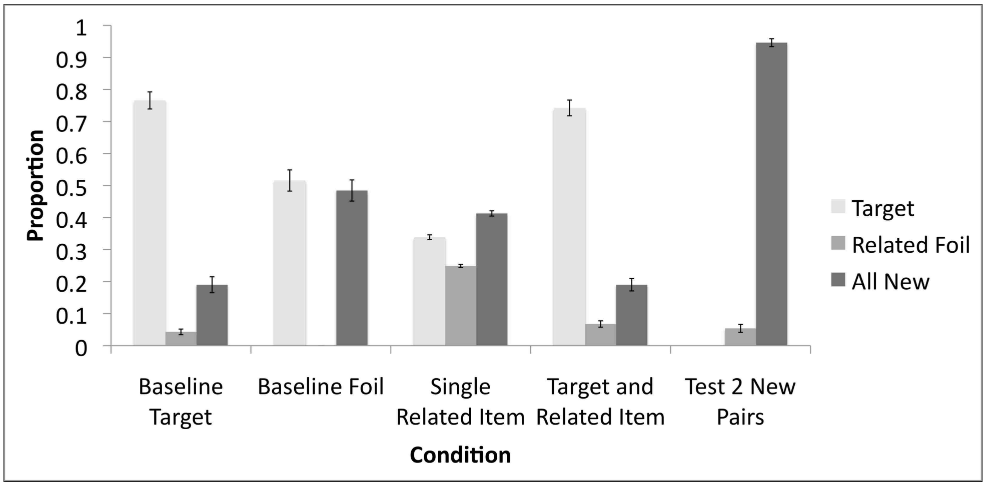

Figure 10. Proportion of Test 2 responses across conditions in Experiment 2.

For the one-way ANOVA examining "all new" responses, Mauchly's test of sphericity was violated, so Greenhouse-Geisser corrected values are reported. There was a significant main effect of all new responding across conditions, F (3.14, 97.28) = 313.50, MSE = 7.39, $p < .001$, $\eta^2 = .910$. Post-hoc comparisons using Bonferroni-corrected p values revealed significant differences between all conditions (ps < .05) except the difference between the *baseline target* and *target and related item* conditions. First, this shows that subjects were able to reject new pairs of items. For the rest of the conditions, the results present a mirror image of the pattern from the above analysis for the *baseline target* and *target and related item* conditions. "All new" responding was highest in the *baseline foil* condition, followed by the *single related item* condition, and lowest in the *baseline target* and *target and related item* conditions. The difference in "all new" responding between the *baseline foil* and *single related item* conditions can be attributed to the influence of the target-related foil on Test 2 responding, as the target-related foil was rarely chosen in the *baseline foil* condition but was a fairly common selection in the *single related item* condition.

To analyze Test 2 responses in relation to Test 1 choices (see Table 2 and Figure 11), I conducted a 2 (Test 1 response: correct rejections vs. false recognition) x 3 (Test 2 response: hit vs. “all new” response vs. related item false alarm) ANOVA. There was a significant main effect of Test 1 response, F (1, 31) = 283.65, MSE = 35.99, $p < .001$, $\eta^2 = .901$, indicating that there were significantly more correct rejections than false recognitions. For the main effect of Test 2 response, Mauchly’s test of sphericity was violated, X^2 (2) = .781, $p = .024$, so Greenhouse-Geisser corrected values are reported. This main effect was significant, F (1.64, 50.85) = 9.37, MSE = 67.15, $p = .001$, $\eta^2 = .232$. “All new” responses were the most prevalent Test 2 response overall, followed by false recognitions, and finally false alarms to the presented foil. For the interaction between Test 1 and Test 2 responding, sphericity was again violated. This interaction was significant, F (1.48, 45.73) = 52.46, MSE = 47.22, $p < .001$, $\eta^2 = .629$. “All new” was the most common response following correct rejections at Test 1, but least common following false recognitions. Meanwhile, false alarms to the target-related foil were the most common response following false recognitions at Test 1, but least common following correct rejections.

	False Recognitions	Correct Rejections	T1 False Alarms
Hits	3.50	20.28	0.59
"All New"	2.50	26.50	0.72
T2 False Alarms	7.34	10.31	0.28

Table 2. Mean number of items correctly recognized (hits), "all new" responses, and false alarms to the related item in Test 2 as a function of responses in the *single related item* condition at Test 1 in Experiment 2. Data from the false recognition and correct rejection columns will be displayed in Figure 11.

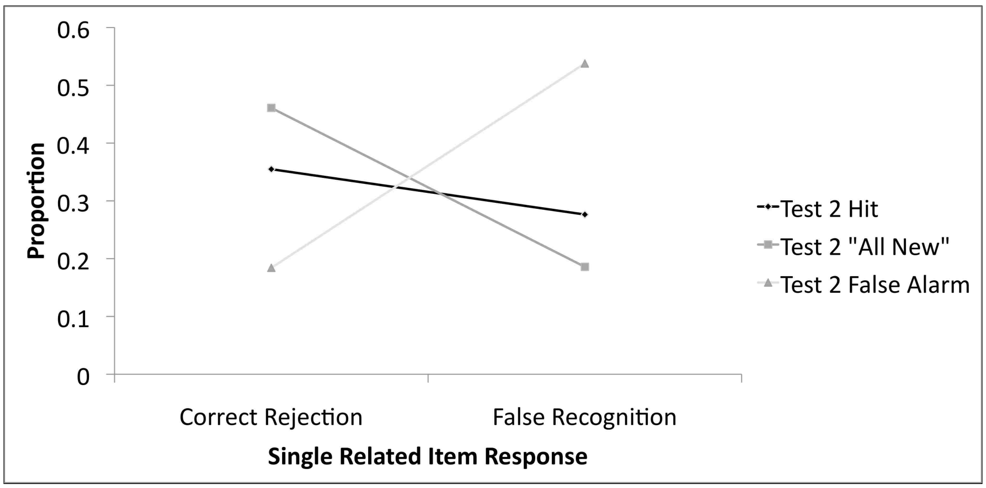

Figure 11. Proportion of Test 2 responding based on responses in *single related item* condition at Test 1 in Experiment 2.

Discussion

Looking first at the Test 1 results, the basic pattern of findings replicated the pattern from Guerin et al. (2012) and Experiment 1.

Looking at corrected recognition at Test 2 based on Test 1 condition, targets assigned to the *baseline target* and *target and related item* conditions were recognized most often at Test 2. It appears that the presence of the target-related foil at Test 2 trials corresponding to the *target and related item* condition did not have much impact on subjects' ability to recognize targets, which aligns with the finding of low foil recognition in the *target and related item* condition in Experiment 1. Targets corresponding to the *baseline foil* and *single related item* conditions exhibited the lowest Test 2 hit rate; however, there were significantly more "all new" responses in the *baseline foil* condition than in the *single related item* condition. This seems to indicate that the presence of the target-related foil in the *single related item* condition has an influence on subjects' subsequent false alarms to the target-related foil, aligning with the results from Experiment 1.

The crucial prediction of this study was that correctly rejecting the foil in comparison to falsely accepting it would lead to subjects mistakenly choosing the target-related foil as the target at Test 2. The results of Experiment 2 provide evidence against this hypothesis. After correctly rejecting the foil in Test 1, subjects' most frequent Test 2 response was to reject all options, i.e., the target and the target-related foil from Test 1. In contrast, after false recognitions subjects most often selected the target-related foil as the target in Test 2. Therefore, it appears that after correct foil rejection, foils are not mistaken as targets. Additionally, while Pilot Study 2 showed that correct rejections in the *single related item* condition lead to activation of perceptual details, false recognitions may only lead to reactivation of gist information. Thus, selection of the target-related foil at Test 1 may cause subjects to use that foil to fill in the missing perceptual details of a concept that they know was encountered previously. Alternatively, it is possible that subjects have better memory for their response at Test 1 than for individual items; therefore, they make responses at Test 2 based on response repetition (as discussed with regard to Experiment 1). This pattern of results is particularly fascinating due to its item specificity; that is, Test 2 responses seem specific to the response given to a particular item at Test 1. Additionally, given the within-subjects design used in these studies, it is clear that these effects manifest themselves within subjects; there is no group manipulation that would allow subjects to habituate to one pattern of correct responses in a condition.

One additional interesting finding is that hit rates following correct rejections and false recognitions at Test 1 are quite similar. This is true despite the significant difference in hits based on Test 1 responding when looking at the raw data. This finding is in contrast to Pilot Study 1, where hits made up a larger proportion of responses following false recognitions than following correct rejections at Test 1. This appears to be due to the presence of the target-related foil among the options at Test 2. It may be that subjects only respond to the target following initial false recognitions in Pilot Study 1 because their preferred choice (the target-related foil) is not presented, and thus they choose the item for which some representation still exists in memory. However, when the target-related foil is presented at Test 2, initial false recognitions lead to subsequent false recognitions of the target-related foil. This is also corroborated by the fact that hit rates following initial correct rejections in both Pilot Study 1 and Experiment 2 were quite similar.

The results of Experiments 1 and 2, taken together, align with the misinformation effect explanation by McCloskey and Zaragoza (1985). Test 1 false recognition more often results in selection of target-related foils as the target when it is presented at a subsequent test, yet the “misinformation” does not appear to eliminate memory for the original target; rather, it may block access to it (see Pilot Study 1). However, in Experiment 2 avoidance of “misinformation” by way of correct rejection at Test 1 does not support target retrieval at Test 2 as in McCloskey and Zaragoza (1985); rather, it often results in perseveration in giving an “all new” response at Test 2, even when such a response is no longer appropriate. This suggests that under some conditions, target retrieval can be blocked in spite of the presence of the target (in opposition to Guerin et al.’s discussion of gist-based false recognition). Future work should examine a second test without the “all new” response option to better understand the impact of response choices on performance.

I have established that correct rejections of a target-related foil tend to result in rejections of both these foils and the associated targets, while false recognitions lead to preferential recognition of that foil when it is presented as an alternative, which, however, does not block access to the target. One question that remains is whether subjects reject items on a perceptual level, or if rejection extends to the conceptual level (i.e., is rejection based on “token” level information or “type” level information). Experiment 3 tested whether subjects retain targets at a conceptual level, particularly following correct rejections of target-related foils.

Experiment 3: ”Type” memory following correct rejections

Experiment 3 was an exploratory study conducted to investigate to what extent perceptual, or “token”, level information and conceptual, or “type”, level information contribute to misses at Test 2, following correct rejections in the *single related item* condition at Test 1. In misinformation studies, researchers attempt to make the misinformation option as plausible as possible. Thus, in misinformation studies different tokens of the same type are used. Hence, it is not possible (and not necessarily relevant) for these studies to examine the “degree” to which memory is lost. However, according to the Fuzzy Trace theory (Brainerd & Reyna, 2002), specific, detailed information about a memory trace decays more quickly than more general information about the trace. This theory posits that gist and perceptual information in memory are stored at different levels and forgotten at differential rates. Meanwhile, the Constructive Memory Framework (Schacter, Norman, & Koutstaal, 1998) suggests that failure to encode specific information in service of encoding broader information may be responsible for effects such as gist-based false recognition. They propose that false recognition is a product of pattern separation failure; that is, the details of an encountered item are not encoded well enough in order to distinguish it from other, similar items, thus leading to confusion when two similar choices are pitted against each other. This theory posits that perceptual information in memories is encoded poorly and is difficult to separate from perceptual information in other similar memories, thus making later retrieval difficult. Taken together, these two theories support the idea that memory processes do not influence gist and perceptual information in identical ways.

In combination with the results of my pilot studies, the above theories inspired me to consider whether Test 2 recognition failures following Test 1 correct rejections stemmed from an inability to recognize the specific target information (“token” level information) or would extend to the concept of the target object itself (“type” level information). Experiment 3 tested whether subjects reject target “types” (e.g., trees in general) or whether they retain the “type” but reject the specific “tokens” that are presented as choices in Test 2. While I did not have specific predictions about the role of conceptual and perceptual information in rejection of targets, I suspected that rejections might have been based on “type” information.

If subjects can still remember the target "type" while "rejecting" the specific target image presented in Test 2, it could mean that the perceptual details of the target were lost after rejection in the first test, and perhaps even that the features of the target-related foil blends with the features of the target in memory. If subjects do not remember seeing any instance of a given item type, it would mean that the rejection affects target memory at a more global level.

This experiment included a verbal label option that allowed subjects to state that they had seen an image of an object, but that neither of the presented choices was the correct target. If subjects do in fact retain conceptual target knowledge after correct rejections, they should choose the verbal label more often than selecting "all new" responses following correct rejections. If subjects are rejecting all "tokens" of a given "type", they should select "all new" responses more often than selecting the verbal label after correct rejections.

Methods

Subjects and Design

Thirty-two undergraduates (21 female, 11 male, mean age = 18.97 years, *SD* = 1.24) were recruited from the Lehigh University participant pool and participated in the study for course credit. Two subjects failed to follow instructions for Test 2, so thirty subjects are included in the analysis. The design was a 4 x 3 within subjects design. Factor 1 was Test 1 condition with 4 levels: *baseline target*, *baseline foil*, *single related item*, *target and related item*. Factor 2 was Test 2 condition with 3 levels: target, new foil, and verbal label.

Materials

The materials were identical to those used in the pilot studies, with the addition of a set of words that represented verbal labels of each set of objects presented during Test 2.

Procedure

Encoding. The encoding phase was identical to that of the previous studies.

Test 1. Test 1 was identical to that of the previous studies.

Test 2. As in the prior studies, Test 2 took place 48 hours after Test 1. For the 144 experimental trials, subjects were presented with a target image, a different "token" of the same "type" as the target that had not been encountered previously, and a verbal label representing the "type" of both images. For the 24 control trials, subjects were presented with two different "tokens" of a "type" that had not been displayed at any prior point in the experiment, along with a verbal label representing the "type" of both images. Subjects were instructed to choose one of the presented images if they remembered seeing that image during encoding. They were instructed to choose the verbal label if they remembered seeing a different "token" of the presented "type" during encoding. Finally, they were given the option to reject all choices as in the previous experiments; subjects were instructed to choose this option when they did not recall seeing any "token" of the presented "type" during the encoding phase. As in the prior experiments, responses were not time-constrained.

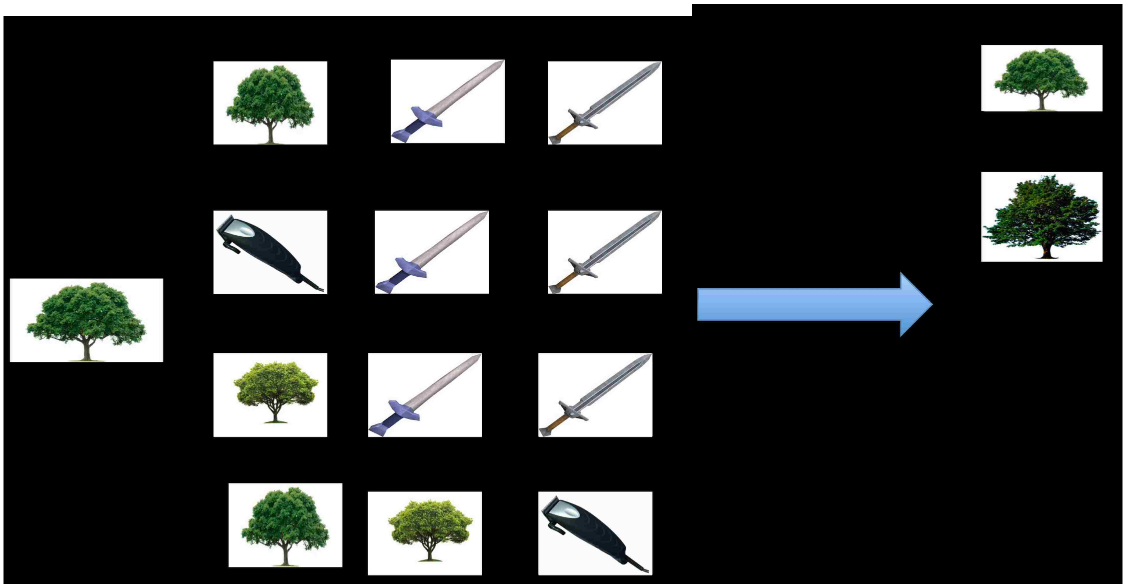

Figure 12. Illustration of conditions in Experiment 3. Items were counterbalanced such that different images were used in all conditions at Test 1.

Results

Test 1

Test 1 results are depicted in Figure 10. Subjects chose the target-related item in the *single related item* condition significantly more often than they chose the standalone item in the *baseline foil* condition, $t(29) = 10.59$, $p < .001$, indicating a standard gist-based false recognition effect. They also committed false recognitions significantly more often in the *single related item* condition than in the *target and related item* condition, $t(29) = 10.51$, $p < .001$. However, false recognition rates in the *target and related item* condition were again significantly higher than false alarm rates to the standalone foil, $t(29) = 1.99$, $p = .056$, and *baseline target* accuracy was again significantly higher than accuracy in the *target and related item* condition, $t(29) = 3.76$, $p = .001$, indicating a reduced but not eliminated gist-based false recognition effect in the *target and related* *item* condition. Once again, this pattern of results replicated Guerin et al. (2012) and the previous work presented here.

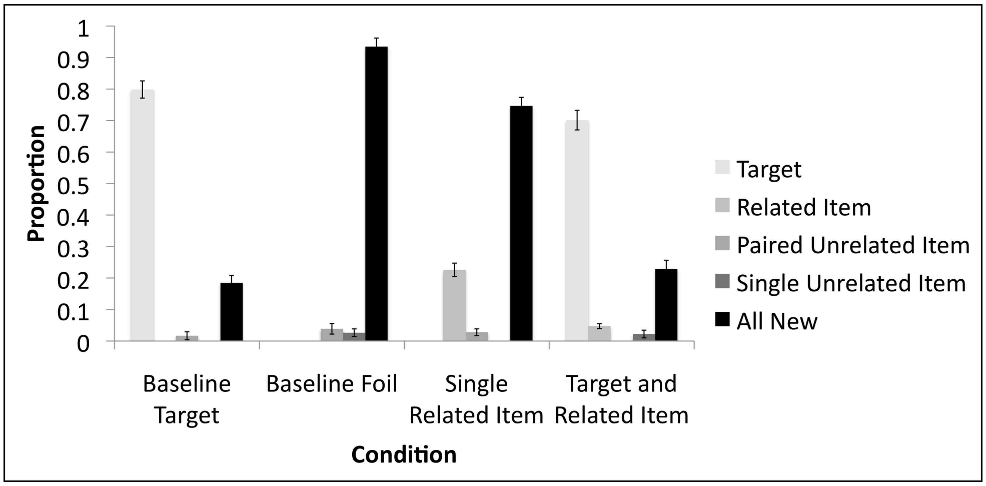

Figure 13. Response proportions at Test 1 across conditions in Experiment 3.

Test 2

I performed three analyses on Test 2 data in order to examine my questions of interest. I first conducted a 4 x 2 repeated-measures ANOVA comparing hits vs. verbal label responses across all four Test 2 conditions that contained target images. This was done in order to directly compare rates of conceptual and perceptual detail memory. I then conducted a one-way repeated measures ANOVA comparing frequencies of "all new" responding across all five Test 2 conditions to know whether misses differed across conditions, and to know whether correct rejections are particularly high when new pairs of images are presented at Test 2. Finally, focusing solely on trials corresponding to the *single related item* condition at Test 1, I ran a 2 x 3 repeated measures ANOVA examining the frequency of Test 2 hits and misses based on whether subjects correctly rejected or falsely recognized target-related foils at Test 1. This will allow for direct comparison of how Test 1 responses influence Test 2 responses in the *single related item* condition while still analyzing the data within the ANOVA framework. All analyses were conducted on raw data, but proportions are depicted in graphs for ease of interpreting the data.

I will first examine the 4 x 2 ANOVA. For the main effect of condition, Mauchly's test of sphericity was violated, so Greenhouse-Geisser corrected values are reported. There was a significant main effect of condition, F (1.87, 54.32) = 83.83, $p <$ .001, $\eta^2 = .743$. Post-hoc comparisons using Bonferroni-corrected p values revealed significant differences between all conditions (all ps < .05) except the *baseline target* and *target and related item* conditions. In the *baseline target* and *target and related item* conditions, hits were much more prevalent responses than verbal label choices, while the difference between responses was much smaller in the *baseline foil* and *single related item* conditions. The main effect of Test 2 response was also significant, F (1, 29) = 69.57, $p < .001$, $\eta^2 = .706$, indicating that there were more hits than verbal label responses overall. Finally, the interaction between condition and Test 2 response was significant, F (3, 87) = 118.04, $p < .001$, $\eta^2 = .803$. Simple main effects analyses revealed significantly more Test 2 hits at the *baseline target* condition (F (1, 29) = 107.05, $p < .001$, $\eta^2 = .787$) and the *target and related item* condition (F (1, 29) = 171.73, $p < .001$, $\eta^2 = .856$). Additionally, there were significant differences in both Test 2 hits (F (3, 27) = 163.44, $p < .001$, $\eta^2 = .948$) and Test 2 verbal label responses (F (3, 27) = 40.21, $p < .001$, $\eta^2 =$ .817) across conditions. It appears that hits were more prevalent in the *baseline target* and *target and related item* conditions as displayed earlier, and that verbal label responses were more prevalent in the *single related item* condition compared to other conditions.

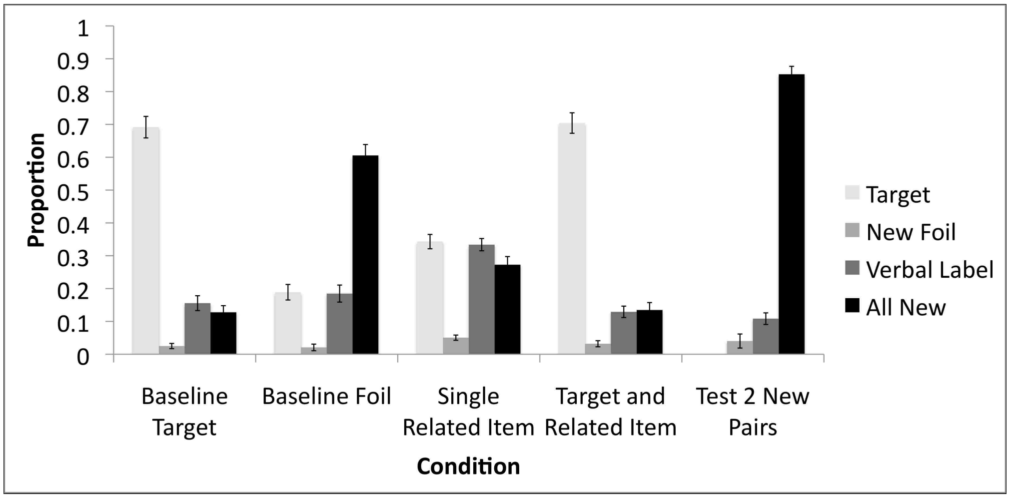

Figure 14. Proportion of Test 2 responses across conditions in Experiment 3.

I then conducted a one-way repeated measures ANOVA to compare the number of "all new" responses across all five conditions. Mauchly's test of sphericity was violated, so Greenhouse-Geisser corrected values are reported. There was a significant main effect of all new responding across conditions, F (2.63, 76.21) = 201.82, MSE = 11.89, $p < .001$, $\eta^2 = .874$. Post-hoc comparisons using Bonferroni-corrected p values revealed significant differences between all conditions (all ps < .05) except the *baseline target* and *target and related item* conditions. The highest levels of "all new" responding appear to come from new pairs of images at Test 2, as well as the *baseline foil* and *single related item* conditions. This indicates that the conditions most likely to produce an "all new" response at Test 1 were also most likely to produce an "all new" response at Test 2.

Finally, I conducted a 2 (Test 1 response, *single related item* condition: correct rejections vs. false recognition) x 3 (Test 2 response: hits to targets vs. "all new" responses vs. verbal label) to examine the effects of false recognition on Test 2 responding (see Table 3 and Figure 15). There was a significant main effect of Test 1 response, F (1, 29) = 124.87, MSE = 54.52, $p < .001$, $\eta^2 = .812$, indicating that there were more correct rejections than false recognitions at Test 1. The main effect of Test 2 response approached significance, F (2, 58) = 2.93, MSE = 48.83, $p = .061$, $\eta^2 = .092$, indicating that there were marginal differences across the possible Test 2 responses. Finally, there was no significant interaction between Test 1 and Test 2 responding (p = .383). In addition, in order to further investigate whether conceptual or perceptual information was the basis of Test 2 rejections, a paired-samples t-test was conducted comparing "all new" responses and verbal label responses following correct rejections. There was no significant difference (p = .432).

	False Recognitions	**Correct Rejections**	**T1 False Alarms**
Hits	6.23	17.77	0.70
"All New"	2.40	16.20	1.03
Verbal Label	6.17	17.73	0.13
T2 False Alarms	1.47	2.03	0.13

Table 3. Mean number of items correctly recognized (hits) vs. "all new" responses in Test 2 as a function of responses in the *single related item* condition at Test 1 in Experiment 3. Hits for this table are defined as correctly responding to exact copies of targets. Data from the false recognition and correct rejection columns will be displayed in Figure 15.

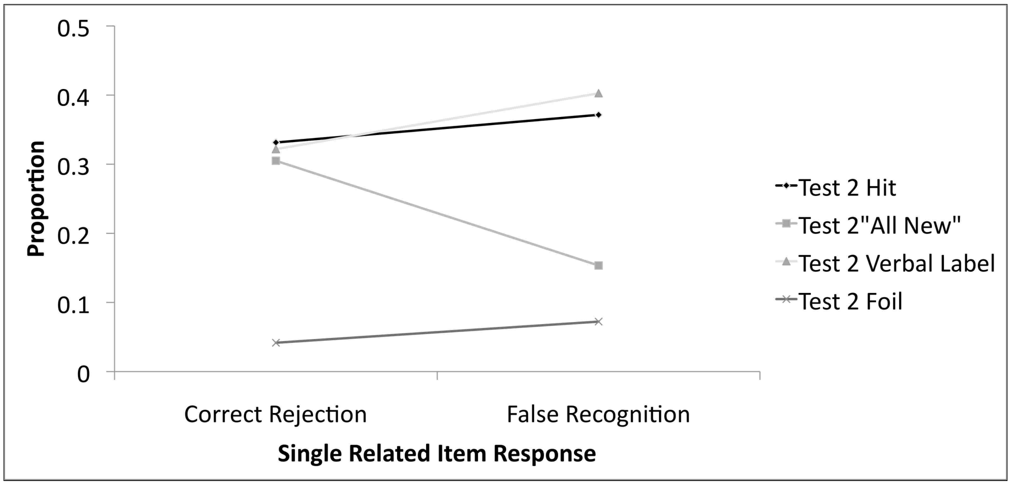

Figure 15. Test 2 responding based on responses in the *single related item* condition at Test 1 in Experiment 3.

Discussion

Looking first at the Test 1 results, the basic pattern of findings replicated the pattern from Guerin et al. (2012) and Experiment 1.

Looking at the patterns of responding at Test 2 based on condition, selection of target images was most common in the *baseline target* and *target and related item* conditions. This is unsurprising, as in both of these conditions the target image had been encountered and was likely selected multiple times, and findings from my previous studies have also found similar results. Meanwhile, verbal label responses were most common in the *single related item* condition. This finding aligns with predictions made by the Constructive Memory Framework regarding the differential memory for perceptual information versus gist information. Encountering a second ”token” of a particular “type” (that is, the target-related foil in the *single related item* condition) may boost “type”-level memory, but could denigrate memory for the encoded target by introducing interference in the form of a second “token” of the same “type”. Conversely, no benefit from Test 1 encounters is given in the *baseline foil* condition, which may explain why both hits and verbal label responses are relatively low in this condition.

Focusing specifically on the *single related item* condition, the patterns of results are very similar following false recognitions and correct rejections. ”All new” responses constitute a larger proportion of the Test 2 responses following Test 1 correct rejections, but this was ultimately insufficient to create a significant interaction. Importantly, “all new” responses and verbal label responses were roughly equal following correct rejections. This brings into question what subjects actually have stored in memory at Test 2, as it appears that, at times, subjects make Test 2 rejections based on perceptual differences between the stored representation and the presented representations. With regard to perceptual details, Brady et al. (2013) found that different perceptual elements fade from memory at differential rates. It might therefore be interesting to manipulate particular perceptual elements in a future study to see if they are differentially affected by the delay. It may also be interesting to investigate what subjects remember from a purely conceptual standpoint (that is, simply ask them what items they remember seeing in lieu of a true Test 2). Alternatively, subjects may view the “all new” response and the verbal label response as roughly equivalent, due to the fact that neither of these responses requires the specific details of the presented images to be retrieved. As a result, subjects may divide their choices between these two options at a chance rate.

It is also possible that subjects sometimes do not engage in veridical retrieval but rather base their responses on their sense of recollection or familiarity with respect to specific items (see Yonelinas, 2002). In this case, the verbal alternative fosters a shift from recollection to familiarity-based recognition. Recollection requires that rich details about a memory are retrieved, while familiarity refers to the broader sense that an item was encountered in the past. In my prior studies, familiarity-based decisions at Test 2 were less likely. In other words, the presence of the verbal label at Test 2 may bias subjects towards selecting the label when they are uncertain about the actual presented response alternatives. Some support for this argument comes from examining correct rejection rates for control trials at Test 2 (that is, trials where no target “types” were presented). Compared to my previous experiments, I found rather low correct rejection rates in this condition. It seems as though the presence of a verbal label causes subjects to evaluate the options less carefully, thus leading to false alarms. In other words, subjects may generally default to the verbal label because it is the response that does not require a clear, categorical yes/no decision. This may attract them to choosing the verbal label as opposed to choosing the images or the “all new” option, which aligns with findings in the decision-making literature (Dhar, 1997). Here consumers use higher selection thresholds when there is a no-choice option, and the no-choice option becomes particularly attractive under difficult conditions. In the case of my study, while subjects are not allowed to opt out of a selection completely, the perceptual ambiguity of the verbal label may present the “safest” choice.

In summary, it appears that rejections at Test 2 are mostly based on rejections of the perceptual details of the presented images. This raises the question of what exactly is stored in memory following an initial correct rejection. Is it possible that only the conceptual information remains, and no perceptual details can be retrieved? Might an entirely new representation with perceptual information from both target and target-related foil have formed? This appears to be an important question for future research.

General Discussion

My thesis investigated the long-term memory consequences that follow both correct rejections and false recognitions in visual recognition memory. In a pilot study, correct rejection in comparison to false recognition of a foil that resembled a previously encountered target image (that is, an item that shared the target "type" but was a different "token") resulted in impaired target recognition 48 hours later. However, a follow-up study revealed that recognition impairments are only found when there is a time delay between Test 1 and Test 2, as no Test 2 difference were observed between correct rejections and false recognitions when the final test occurred immediately after the initial test. This suggests that subjects reject foils at Test 1 because they successfully access and compare the foils to memory for original targets, and that recognition impairment evolves over time. The series of studies in this thesis examined potential explanations for these delayed recognition impairments following initial correct rejections. Experiment 1 tested the hypothesis that correct rejections at Test 1 result in richer encoding of the target-related foil. This might occur because rejection of the target-related foil is difficult; the subject must directly compare the details of the foil to the stored target representation in order to correctly reject. This might then lead the target-related foil to become encoded more elaborately (Craik & Lockhart, 1972). Experiment 2 tested the related hypothesis that correct rejection of the target-related foil at Test 1 may lead to a false assumption that the foil is the target at Test 2. This was hypothesized based on the idea that Test 1 correct rejections require target activation, thus making the representation of the target susceptible to change (as in prior work on reconsolidation, e.g. Hupbach et al., 2007).

The patterns of results from these experiments do not support these hypotheses. Rather, it appears that correct rejections often lead to later rejections of both targets and target-related foils, while Test 1 false recognitions frequently cause subjects to mistake the target-related foil they chose at Test 1 for the target at Test 2. This may be due to the fact that target-related foils were encountered more recently than targets, which increases the fluency with which foils are processed and causes subjects to falsely believe they are the actual targets (Whittlesea & Williams, 2001). Experiment 3 explored to whether conceptual level information was retained at Test 2 following correct rejections at Test 1. The results of this experiment are less clear, as it seems that there is evidence for rejections on both the "type" and "token" level.

I will address this complex set of results from a variety of perspectives, some of which were mentioned in the Introduction and some of which are new accounts that may provide explanations for the patterns of data seen in these studies. I will first address retrieval-based accounts of the results; that is, explanations based on various retrieval failures and how this work may also inform work in these fields of research. Then, I will address storage-based accounts; that is, explanations based on the possibility that stored representations in memory are altered. These storage-based accounts must be addressed carefully, as it is difficult to directly test change in memory storage. However, the literature in this section has provided some evidence for storage change and has fascinating implications for both the current work and future directions.

Retrieval-Based Explanations

Although not designed to explicitly misinform participants about previously encountered information, the present experiments can nevertheless inform the misinformation effect literature that was discussed extensively in the Introduction. While Loftus (1977) argued that presentation of misinformation following a witnessed event could alter the details of the original memory trace such that the misinformed details were now integrated into the original memory, McCloskey and Zaragoza (1985) provided evidence that the paradigm used by Loftus was insufficient to make claims about the fate of the original memory. They claimed that because not all subjects in Loftus' work would remember the original information, those subjects would be biased to choose the misinformation, and still other subjects that did in fact remember the original information may still choose the misinformation due to source monitoring error (see Johnson, Hashtroudi, & Lindsay, 1989). McCloskey and Zaragoza showed that when the original information was pitted against new information, the presentation of post-event misinformation did not lead to memory impairment. This suggested that task demands were responsible for Loftus' results as opposed to changes in the original memory trace, and a variety of alternative tests and conditions have been used to attempt to settle this debate (e.g., Zaragoza, McCloskey, & Jamis, 1987; Belli, 1989; Tversky & Tuchin, 1989).

Considering the results of the present experiments, there seems to be a clear link between the pattern of findings following false recognitions at Test 1 and McCloskey and Zaragoza (1985)'s findings. In some sense, when subjects commit a false recognition at Test 1 they have fallen for the "misinformation". Experiment 1, in which subjects were asked to select previous items independent of their source (that is, encoding and Test 1) showed that subjects were more likely to correctly identify the target-related foil at Test 2 after Test 1 false recognitions than after Test 1 correct rejections. In Experiment 2, foil selection was the incorrect choice, because subjects were asked to select encoded target items only. However, after committing a false recognition at Test 1, subjects still frequently selected foils at Test 2. This aligns with both Loftus' (1977) and McCloskey and Zaragoza's findings that when a misleading detail is presented in a later memory test, subjects often choose the misleading detail as part of the original narrative. However, Pilot Study 1 also showed that subjects more often have the ability to retrieve the target at Test 2 when the target-related foil is absent, which corroborates the role of response bias that McCloskey and Zaragoza showed in their studies. However, to my knowledge, no prior work on the misinformation effect has integrated the possibility that subjects can decide that both choices are incorrect, which makes the pattern of results following Test 1 correct rejections difficult to reconcile with this literature. In the studies presented here, "all new" responding does not map onto "no misinformation" control conditions in misinformation effect studies, because those studies do not expose subjects to any misinformation at all. Conversely, in this thesis, subjects are always exposed to misleading "tokens", but are sometimes able to identify them as misleading and reject them. Future work on the misinformation effect should investigate whether the ability to overtly reject misinformation influences subsequent retrieval of the original episode.

The results of these experiments also may help understand what role retrieval failures actually play in gist-based false recognition. Recall that Guerin et al. (2012) showed that false recognition is reduced when a target is part of the Test 1 triplet. Based on this, they argue that retrieval failure is the primary contributor to this effect; people have already stored all of the requisite details in order to choose correctly, but do not properly access and use them unless the target itself is present. However, the studies presented here seem to suggest that, with a longer delay, even the presence of the target is often not enough to promote recognition, especially following an initial correct rejection of a target-related foil. This seems to imply the possibility that what is stored in memory is changed over the course of the delay, or fades to the degree that retrieval is not possible even with the presence of the target, which would contradict Guerin et al.'s theory with regard to the special role of retrieval mechanisms in gist-based false recognition. These storage changes might occur within the context of the reconsolidation framework discussed in the Introduction. Given that the target representation is activated in order to make correct rejections in the *single related item* condition at Test 1 (per Pilot Study 2), the target representation might become susceptible to alteration by the presence of the target-related foil (see the next section for further, more detailed discussion of the potential effects of reconsolidation).

One alternative mechanism for the supposed retrieval failure seen in these studies is that the response given to an image at Test 1 may influence the subjective "truth" of that image; that is, Test 1 selections may influence subjects' beliefs about what was encoded. Previous work on the effects of multiple-choice testing has shown that an initial multiple-choice test is beneficial for performance on a subsequent multiple-choice test (Hogan & Kintsch, 1971), but when a lure is chosen on an initial multiple-choice test, the same lure is often produced on a later cued-recall test (Roediger & Marsh, 2005). This seems to indicate that a piece of information and the response given to that piece of information are closely tied, regardless of whether the initial response was correct or not.

Considering this work in the framework of the results presented here, subjects may perseverate in an answer given at Test 1 on a later memory test, which may explain why subjects so often repeat their Test 1 responses when the Test 2 options give them the option to do so. Additionally, the studies presented here suggest that response perseveration may continue even with changing task demands from Test 1 to Test 2 as seen in Experiment 1. When subjects correctly responded “all new” to a target-related foil at Test 1, they often failed to choose that foil at Test 2 when task demands called for them to do so. This indicates that previous selection of an encountered item may be a powerful influence over later selection of that item.

Another possible explanation for the pattern of data in these studies is the possibility that subjects engage in a sort of directed forgetting of items that were rejected at Test 1. In directed forgetting studies, subjects typically are asked to encode a series of items either one at a time (item-method) or as an entire list (list-method). In the item-method paradigm, subjects are given explicit cues to forget or remember each particular item, and typically it is found that to-be-remembered items are recalled better than to-be-forgotten items. This effect is often attributed to active rehearsal; subjects continually rehearse an item when they receive a cue to remember the item, while they are able to terminate rehearsal when they receive a cue to forget the item (Basden et al., 1993). In the list-method paradigm, subjects are typically asked to forget one list and remember another; again subjects typically remember items from the to-be-remembered list compared to the to-be-forgotten list. This effect has often been explained by inhibition accounts that suggest that the to-be-forgotten list is inhibited once the cue is given and thus is more difficult to recall (Geiselman, Bjork, & Fishman, 1983; although see also

Sahakyan & Kelley, 2002 for an alternative account). In all of these studies, words are typically the stimuli of choice due to their sensitivity to rehearsal and inhibition effects. However, despite the use of words as the stimuli of choice in these studies due to their sensitivity to rehearsal effects in particular, Hourihan, Ozubko, and Macleod (2009) showed that visual images can also undergo forms of rehearsal that lead to directed forgetting effects, even when no verbal label can be attached to that image. While to my knowledge no directed forgetting studies have directly manipulated particular perceptual details, Hourihan et al.'s work suggests that it may be possible to engage in rehearsal of specific perceptual details and perhaps find directed forgetting effects tied specifically to visual perception.

Considering directed forgetting in the context of the experiments presented here is difficult, but may provide some interesting avenues for future research. It seems that, across studies, "all new" responses at Test 1 engage subjects in a form of item-method directed forgetting where they lose the ability to recognize rejected items later. However, if directed forgetting can indeed qualify as an explanation for these results, the specific form of directed forgetting that takes place during these studies would have some unique qualities not explained by the current literature. First, the directed forgetting in this study is self-motivated, whereas prior work used external cues to motivate subjects to remember or forget. Additionally, in prior directed forgetting work, the cue to remember or forget an item or list is given immediately after encoding, while these studies indicate that directed forgetting cues may be effective well after encoding and potentially with a great deal of intervening information processed between the encoding phase and the "cue". Future directed forgetting work should investigate the effects of delay between the encoding of items and the remember/forget cue, as well as explore ways of testing whether self-generated cues might be effective in selectively forgetting encoded information.

Storage-Based Explanations

One additional possibility with regard to potential changes in target representations is that the encoded target and the target-related foil encountered during Test 1 sometimes morph together in some way, particularly following correct rejection of target-related foils. This seems most likely to happen after correct rejections given that activating the target representation is required in order to give a reject response (per Pilot Study 2). The possibility of "morphs" like these is not unprecedented in the memory literature. Weinberg, Wadsworth, and Baron (1983) found in an investigation of the misinformation effect that subjects who saw a yellow yield sign in a series of slides and were given false information about the presence of a stop sign were more likely to report having seen a red yield sign than a yellow yield sign compared to those who did not receive misinformation, indicating that the perceptual details of the stop sign had blended with existing memory for a yield sign to create a morphed representation in memory. Subsequent work on memory blends has also found that prior knowledge about certain representations can also factor into memory blends (Belli, 1988), indicating that a wide variety of factors including past experiences can cause details in memory to become altered. This is particularly relevant for the work presented here; the stimuli presented are very common objects of which subjects can be expected to have some existing representation. Thus, it is possible that existing knowledge about the presented images may impact the way presented representations are encoded into memory. In a related vein, hypotheses of “partial degradation” with regard to the misinformation effect have been offered in past work (e.g. Johnson & Raye, 1981, Chandler, 1989, Tversky & Tuchin, 1989). These theories state that memories have a variety of features, and each feature of a memory is not necessarily stored at one uniform strength. Rather, features may differ in their strength, which could result in blended memories in the presence of interfering information (see also Brady et al., 2013, for evidence that perceptual elements of visual memories are forgotten at different rates).

With regard to the present work, it is possible that correct rejection of target-related foils at Test 1 support a blended representation in memory, thus leading to the persistent Test 2 rejections seen across these studies. The reconsolidation account may be able to explain these results. In Pilot Study 2, it was shown that target reactivation is required to correctly reject in the *single related item* condition. Therefore, if the target is reactivated at for *single related item* trials at Test 1, thus leading to correct rejections, details from the presented image may influence the stored representation of the target. This aligns with previous findings in the reconsolidation literature of both memory impairment (e.g. Forcato, Rodriguez, & Pedeira, 2011) and memory change in the presence of new information (e.g. Hupbach et al., 2007). In Experiments 1 and 2, as well as Pilot Study 1, a correct rejection at Test 1 in the *single related item* condition followed by a 48-hour delay led to subsequent failure to recognize encoded targets and as well as the target-related foils that had been rejected. In Experiment 3, subjects were given the option of choosing a verbal label over any of the image choices, and the presence of this option eliminated the differences in Test 2 memory based on Test 1 *single related item* responding that had been seen in the other experiments. This pattern provides evidence in favor of the possibility that a perceptual representation lying somewhere between target and target-related foil is what is present in memory at Test 2 following correct rejection of target-related foils. However, there is still conflict between these results and previous work on memory blends. Pointing back to the study by Weinberg et al. (1983), misinformed subjects were given conflicting details in a misleading narrative, thus leading to the apparent integration of the details from both sources into memory. However, in that study subjects who were not in the misinformation condition did not encounter misinformation at all, while in the studies presented here all subjects encounter a misleading piece of information (the target-related foil). This brings into question why, if a representational blend were a possibility, only correct rejections at Test 1 seem to foster evidence of these blends. It is possible that subjects arrive at a verbal label response via different routes based on whether they correctly rejected or falsely recognized a target-related foil at Test 1. After correct rejections, subjects may blend the details of the original target representation and the target-related foil as a result of the reactivation of the target representation on correct rejection trials, while after false recognitions, subjects may simply believe that the target-related foil representation is the target, and thus choose the verbal label because they cannot identify the particular "tokens" presented. In order to test the idea that the perceptual details of target and target-related foil may blend together, it could be of interest to create blends of the perceptual details between targets and target-related foils to see whether these might match what subjects have stored in memory. Likewise, it might be useful to offer subjects a verbal label only at Test 2 and ask them to describe the details of the item they have represented in memory for this label.

Additionally, trace decay and interference-based explanations for forgetting briefly discussed in the Introduction may help explain these results. Hardt et al. (2013) posited that trace decay is a more likely explanation for forgetting when fully consolidated memories are lost, while interference is a more likely explanation when memories that are either not fully consolidated or have been reactivated through the retrieval process are lost. Applying this framework to the results of these studies, the pattern of results at Test 2 following false recognitions might be explained through an interference account. Perhaps a Test 1 false recognition, made in an attempt at accurate target retrieval, causes the details of the target-related foil to interfere with the details of the target, thus leading subjects to believe that the representation of the foil is actually the target representation. However, interference effects cannot explain Test 1 correct rejections, as I have shown in Pilot Study 2 that these rejections are based on accurate target retrieval. Based on this framework, then, the forgetting that is shown at Test 2 following Test 1 correct rejections must have a trace decay basis. But why are memories following correct rejections specifically susceptible

Another possible mechanism by which the original target memory may be altered, and a mechanism that is implicated by Hardt et al. (2013) as having a crucial role in trace decay, is the process of memory tagging and selective consolidation during sleep. Recent evidence suggests that memories can be given explicit or implicit tags that may enhance or reduce their candidacy for preferred consolidation that takes place during sleep (for a recent review, see Stickgold & Walker, 2013). For example, being informed about a forthcoming test on newly learned material selectively enhances later retrieval of that material after sleep compared to an equivalent amount of time awake (Wilhelm et al., 2011). Directed forgetting effects also increase when subjects have slept in the intervening period between learning and test (Saletin et al., 2011), providing an indication that during sleep memories tagged as relevant (that is, to be remembered) are consolidated at the expense of memories tagged as irrelevant (to be forgotten). However, no differences were found between sleep and awake groups when subjects were asked to think about or suppress individual items in a paired-associates learning task (Fischer, Diekelmann, & Born, 2011). This confirmed that the future relevance of information is the driving mechanism behind selective consolidation, as the task used by Fischer et al. did not ask subjects to retain specific learning items across the sleep period. Additionally, this learning can spread beyond the particular items chosen for selective consolidation. For example, subjects who were given the remote-associates task (Mednick, 1962) and were retested on incorrect responses following a nap performed significantly better on retest than subjects who did not sleep (Cai et al., 2009). Payne et al. (2009) also showed that subjects are more likely to recall critical lures in the Deese-Roediger-McDermott false memory paradigm (Roediger & McDermott, 1995) after both a full night of sleep and a shorter nap compared to an equivalent amount of time awake. This body of work suggests a powerful influence of sleep on memory consolidation processes.

How does this framework map onto the present findings? Consider what a correct rejection or a false recognition in the *single related item* condition represents. When committing a false recognition, one overtly declares the relevance of a particular item (the target-related foil). Thus, when the target-related foil is presented as an option at Test

2, which takes place after a period during which subjects have presumably slept, the foil has been selectively consolidated, making it a preferred option when presented as an alternative at Test 2. The correct and incorrect responding to the foil seen in Experiments 1 and 2, respectively, evidences this. However, this effect does not present itself in Pilot Study 2 when there was no delay between the two tests. This does not necessarily impact memory for the original target, and it does not preclude the possibility that subjects can still access that target if it is the most viable option at Test 2 (as in Pilot Study 1). Conversely, when a subject correctly rejects the items in the *single related item* condition (including the target-related foil), all items in that trial have been declared irrelevant. Therefore, the target-related foil is tagged as an item that does not need to be retained, and it might be lost from memory in the sleep consolidation process. This may explain the poor recognition of the target-related foil after correct rejections observed in Experiment 1. It is furthermore possible that this process can cause the target to be forgotten as well, as Cai et al. (2009) showed using the remote-associates task. Recall that in Pilot Study 2, I established that subjects correctly reject the target-related foil in the *single related item* condition at Test 1 as a result of the target's continued presence in memory. It follows, then, that the target must be activated to some degree in order for a correct rejection to occur. If indeed rejection of the target-related foil takes place while the target representation is active, then both foil and target may be subjected to a decreased level of consolidation. This could be a viable explanation for why correct rejections seem to lead to a failure to recognize the original target as seen in Pilot Study 1 and Experiment 2. Additional research should examine the possible influence of sleep and relevance tags in memory, both using this paradigm specifically and in visual recognition memory more broadly.

Conclusions

Overall, the studies of my Master's thesis provide a clear pattern for how responding at an initial memory test modulates subsequent long-term memory. Following a false recognition at Test 1, subjects are most likely to select at Test 2 the same foil they chose at Test 1 if it is offered as an option. However, Pilot Study 1 also showed that when the foil chosen at Test 1 is not an option at Test 2, subjects have the ability to recover and select the target. Conversely, following a correct rejection at Test 1, subjects are most likely to reject all presented options at Test 2. This remains true when for both encoded targets and foils presented at Test 1; however, while subjects might reject specific "tokens" it appears that they do not reject "type" information.

These results raise some important directions for future work. Future studies should further examine the conditions under which people may and may not be able to recover from initial memory error. Furthermore, future work should further investigate the possibility that stored representations in memory may be altered. This may be done with the use of subjective reports where subjects indicate directly the perceptual details of images they believe they have seen. More research should also be done to investigate the role of sleep in these effects, particularly whether or not selection of an item can be interpreted as a relevance tag, and whether or not rejection of a set of items can cause those items to be deemed irrelevant. Finally, the studies presented here should be extended to more practical applications such as facial recognition. This kind of work could have important implications for such situations as attempting to identify a criminal from a police

References

Anderson, R., & Pichert, J. (1978). Recall of previously unrecallable information following a shift in perspective. *Journal of Verbal Learning and Verbal Behavior*, *17*(1), 1-12.

Baddeley, A. D., & Scott, D. (1971). Short term forgetting in the absence of proactive interference. *The Quarterly Journal of Experimental Psychology, 23*(3), 275-283.

Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology.* New York: Macmillan.

Basden, B. H., Basden, D. R., & Gargano, G. J. (1993). Directed forgetting in implicit and explicit memory tests: A comparison of methods. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*(3), 603-616.

Belli, R. F. (1988). Color blend retrievals: Compromise memories or deliberate compromise responses? *Memory & Cognition, 16*(4), 314-326.

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, *105*(38), 14325-14329.

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2013). Real-world objects are not represented as bound units: Independent forgetting of different object details from visual memory. *Journal of Experimental Psychology: General, 142*(3), 791-808.

Brainerd, C. J., & Reyna, V. F. (2002). Fuzzy-trace theory and false memory. *Current Directions in Psychological Science, 11*, 164–169.

Bransford, J. D., & Franks, J. J. (1971). The abstraction of linguistic ideas. *Cognitive Psychology*, *2*, 331-350.

Cai, D. J., Mednick, S. A., Harrison, E. M., Kanady, J. C., & Mednick, S. C. (2009). REM, not incubation, improves creativity by priming associative networks. *Proceedings of the National Academy of Sciences, 106*(25), 10130-10134.

Chan, J. C. K., Thomas, A. K., & Bulevich, J. B. (2009). Recalling a witnessed event increases eyewitness suggestibility: The reversed testing effect. *Psychological Science, 20*(1), 66-73.

Chandler, C. C. (1989). Specific retroactive interference in modified recognition tests: Evidence for an unknown cause of interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 256-265.

Cowan, N., Beschin, N., & Della Sala, S. (2004). Verbal recall in amnesiacs under conditions of diminished retroactive interference. *Brain, 127*(4), 825-834.

Craik, F. I., & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior, 11*(6), 671-684.

Dhar, R. (1997). Context and task effects on choice deferral. *Marketing Letters, 8*(1), 119-130.

Fischer, S., Diekelmann, S., & Born, J. (2011). Sleep's role in the processing of unwanted memories. *Journal of Sleep Research, 20*(2), 267-274.

Forcato, C., Rodríguez, M. L., & Pedreira, M. E. (2011). Repeated labilization-reconsolidation processes strengthen declarative memory in humans. *PLoS One, 6*(8).

Geiselman, R. E., Bjork, R. A., & Fishman, D. L. (1983). Disrupted retrieval in directed forgetting: a link with posthypnotic amnesia. *Journal of Experimental Psychology: General, 112*(1), 58-72.

Guerin, S. A., Robbins, C. A., Gilmore, A. W., & Schacter, D. L. (2012). Retrieval failure contributes to gist-based false recognition. *Journal of Memory and Language*, *66*(1), 68-78.

Hansen, M. (2001). Scoping out eyewitness IDs. *ABAJ*, *87*, 39-39.

Hardt, O., Einarsson, E. O., & Nader, K. (2010). A bridge over troubled water: Reconsolidation as a link between cognitive and neuroscientific memory research traditions. *Annual Review of Psychology, 61*, 141-167.

Hardt, O., Nader, K., & Nadel, L. (2013). Decay happens: The role of active forgetting in memory. *Trends in Cognitive Sciences, 17*(3), 111-120.

Hauswald, A., & Kissler, J. (2008). Directed forgetting of complex pictures in an item method paradigm. *Memory, 16*(8), 797-809.

Hogan, R. M., & Kintsch, W. (1971). Differential effects of study and test trials on long-term recognition and recall. *Journal of Verbal Learning & Verbal Behavior, 10*, 562-567.

Hourihan, K. L., Ozubko, J. D., & MacLeod, C. M. (2009). Directed forgetting of visual symbols: Evidence for nonverbal selective rehearsal. *Memory & Cognition, 37*(8), 1059-1068.

Hupbach, A., Gomez, R., Hardt, O., & Nadel, L. (2007). Reconsolidation of episodic memories: A subtle reminder triggers integration of new information. *Learning & Memory*, *14*(1-2), 47-53.

Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin, 114*(1), 3-28.

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review, 88*(1), 67-85.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition, 27*(2), 117-143.

Koutstaal, W., & Schacter, D. L. (1997). Gist-based false recognition of pictures in older and younger adults. *Journal of Memory and Language*, *37*(4), 555-583.

Lindsay, D. S. (1990). Misleading suggestions can impair eyewitnesses' ability to remember event details. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 16*, 1077-1083.

Lindsay, D. S., & Johnson, M. K. (1989). The eyewitness suggestibility effect and memory for source. *Memory & Cognition, 17*, 349-358.

Loftus, E. F. (1975). Leading questions and the eyewitness report. *Cognitive Psychology, 7*(4), 560-572.

Loftus, E. F. (1977). Shifting human color memory. *Memory & Cognition*, *5*(6), 696-699.

Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory, 12*(4), 361-366.

Loftus, E. F., Miller, D. G., & Burns, H. J. (1978). Semantic integration of verbal information into a visual memory. *Journal of Experimental Psychology: Human Learning & Memory*, *4*(1), 19-31.

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior, 13*(5), 585-589.
Logan, G. D. (2002). An instance theory of attention and memory. *Psychological review, 109*(2), 376.
Lyle, K. B., & Johnson, M. K. (2006). Importing perceived features into false memories. *Memory*, *14*(2), 197-213.
Macleod, C. M. (1999). The item and list methods of directed forgetting: Test differences and the role of demand characteristics. *Psychonomic Bulletin & Review, 6*(1), 123-129.
Mazzoni, G., & Memon, A. (2003). Imagination can create false autobiographical memories. *Psychological Science*, *14*(2), 186-188.
McCloskey, M., & Zaragoza, M. (1985). Misleading postevent information and memory for events: Arguments and evidence against memory impairment hypotheses. *Journal of Experimental Psychology: General*, *114*(1), 1-16.
Mednick, S. (1962). The associative basis of the creative process. *Psychological Review, 69*(3), 220.
Nader, K., Schafe, G. E., & Le Doux, J. E. (2000). Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature, 406*(6797), 722-726.
Payne, J. D., Schacter, D. L., Propper, R. E., Huang, L. W., Wamsley, E. J., Tucker, M. A., ... & Stickgold, R. (2009). The role of sleep in false memory formation. *Neurobiology of Learning and Memory, 92*(3), 327-334.
Roediger, H. L., & Marsh, E. J. (2005). The positive and negative consequences of multiple-choice testing. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*(5), 1155-1159.
Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(4), 803-814.
Rosenzweig, E. S., Barnes, C. A., & McNaughton, B. L. (2002). Making room for new memories. *Nature Neuroscience, 5*(1), 6-8.
Sahakyan, L., & Kelley, C. M. (2002). A contextual change account of the directed forgetting effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*(6), 1064-1072.
Saletin, J. M., Goldstein, A. N., & Walker, M. P. (2011). The role of sleep in directed forgetting and remembering of human memories. *Cerebral Cortex*, 34-41.
Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology, 49*, 289–318.
Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior, 6,* 156-163.
Skaggs, E.B. (1933). A discussion on the temporal point of interpolation and degree of retroactive inhibition. *Journal of Computational Psychology, 16*, 411–414
Standing, L. (1973). Learning 10000 pictures. *Quarterly Journal of Experimental Psychology, 25*(2), 207-222.

Stickgold, R., & Walker, M. P. (2013). Sleep-dependent memory triage: Evolving generalization through selective processing. *Nature Neuroscience, 16*(2), 139-145.

Thorndike, E.L. (1913). Educational Psychology. Teachers College, Columbia University

Tversky, B., & Tuchin, M. (1989). A reconciliation of the evidence on eyewitness testimony: Comments on McCloskey and Zaragoza. *Journal of Experimental Psychology: General*, *118*(1), 86-91.

Uncapher, M.R., and Rugg, M.D. (2009). Selecting for memory? The influence of selective attention on the mnemonic binding of contextual information. *The Journal of Neuroscience*, *29,* 8270–8279.

Uncapher, M.R., Hutchinson, J.B., and Wagner, A.D. (2011). Dissociable effects of top-down and bottom-up attention during episodic encoding. *The Journal of Neuroscience, 31*, 12613–12628.

Weinberg, H. I., Wadsworth, J., & Baron, R. S. (1983). Demand and the impact of leading questions on eyewitness testimony. *Memory & Cognition*, *11*(1), 101-104.

Whittlesea, B. W., & Williams, L. D. (2001). The discrepancy-attribution hypothesis: I. The heuristic basis of feelings and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27*(1), 3-13.

Wilhelm, I., Diekelmann, S., Molzow, I., Ayoub, A., Mölle, M., & Born, J. (2011). Sleep selectively enhances memory expected to be of future relevance. *The Journal of Neuroscience, 31*(5), 1563-1569.

Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language, 46*(3), 441-517.

Zaragoza, M. S., McCloskey, M., & Jamis, M. (1987). Misleading postevent information and recall of the original event: Further evidence against the memory impairment hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*(1), 36-44.

Matthew Sabia
sabiam3@gmail.com
570-903-5003

Department of Psychology
Lehigh University
17 Memorial Dr E
Bethlehem, PA 18015

Education

Pursuing Ph.D., Psychology, Lehigh University, Bethlehem, PA
Concentration: Cognitive Psychology

B.S. Psychology, The University of Scranton, Scranton, PA
May 2013

General Research Interests

- Visual memory
- Visual cognition
- Memory reconsolidation

Research Projects

2014-2015 Master's Thesis
Department of Psychology, Lehigh University
Exploring the long-term consequences of false and correct recognition

2014 First Year Research Project
Department of Psychology, Lehigh University
Long-term consequences of gist-based false recognitions

2013 Primary Investigator: Bryan Burnham
Department of Psychology, The University of Scranton
Contingent attentional capture by items in selectively ignored locations

2011-2012 Primary Investigator: Bryan Burnham
Department of Psychology, The University of Scranton
Components of working memory and visual selective attention

Professional Organizations

Eastern Psychological Association, Member
Association for Psychological Science, Graduate Student Member

Publications

Burnham, B. R., Sabia, M., & Langan, C. (2014). Components of working memory and visual selective attention. *Journal of Experimental Psychology: Human Perception and Performance, 40*(1), 391-403.

Poster Presentations

Sabia, M. T., & Hupbach, A. (2015, May). The puzzle of repeated recognition: Why do correct foil rejections sometimes lead to subsequent target misses? *Poster presented at the 27th annual meeting of the Association for Psychological Science, New York, NY.*

Sabia, M., & Hupbach, A. (2014, November). When correct rejections lead to memory errors. *Poster presented at the 55th annual meeting of the Psychonomic Society, Long Beach, CA.*

Langan, C. G., Sabia, M. T., & Burnham, B. R. (2012, March), Memory load and control of visual attention. *Poster presented at the 83rd annual meeting of the Eastern Psychological Association, Pittsburgh, PA.*

Teaching Experience

2014	Lab Instructor – Statistical Analysis of Behavioral Data Lab, Lehigh University
2014	Graduate Teaching Assistant – Statistical Analysis of Behavioral Data, Lehigh University
2014-2015	Lab Instructor – Mind and Brain Lab, Lehigh University
2014-2015	Graduate Teaching Assistant – Mind and Brain, Lehigh University
2013	Graduate Teaching Assistant – Cognitive Psychology, Lehigh University
2013	Undergraduate Teaching Assistant – Clinical Psychology, The University of Scranton
2012	Undergraduate Teaching Assistant – Research Methods Lab, The University of Scranton

Service

2015	Chaired “Cognitive Psychology: Perception and Memory” session at Lehigh Valley Associate of Independent Colleges (LVAIC) Conference
2014-2015	Served as graduate student member of the Graduate Committee at Lehigh University

Honors and Awards

2012	Psi Chi, the International Honor Society in Psychology