Retrieval Failure Contributes to Gist-Based False Recognition

Citation	Guerin, Scott A., Clifford A. Robbins, Adrian W. Gilmore, and Daniel L. Schacter. 2012. Retrieval failure contributes to gist-based false recognition. Journal of Memory and Language 66(1): 68-78.
Published Version	doi:10.1016/j.jml.2011.07.002
Accessed	February 19, 2015 12:03:11 PM EST
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:10718357
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

(Article begins on next page)

Retrieval Failure Contributes to Gist-Based False Recognition

Scott A. Guerin[1], Clifford A. Robbins[1], Adrian W. Gilmore[2], Daniel L. Schacter[1]

[1] Department of Psychology,

Harvard University

[2] Department of Psychology

Washington University, St. Louis, MO

Corresponding Author:
Scott A. Guerin
33 Kirkland Street
Cambridge, MA 02138
guerin@wjh.harvard.edu
tel: 617-496-5909
fax: 617-496-3122

ABSTRACT

People often falsely recognize items that are similar to previously encountered items. This robust memory error is referred to as *gist-based false recognition*. A widely held view is that this error occurs because the details fade rapidly from our memory. Contrary to this view, an initial experiment revealed that, following the same encoding conditions that produce high rates of gist-based false recognition, participants overwhelmingly chose the correct target rather than its related foil when given the option to do so. A second experiment showed that this result is due to increased access to stored details provided by reinstatement of the originally encoded photograph, rather than to increased attention to the details. Collectively, these results suggest that details needed for accurate recognition are, to a large extent, still stored in memory and that a critical factor determining whether false recognition will occur is whether these details can be accessed during retrieval.

INTRODUCTION

Human memory is not a literal reproduction of the past, like a photograph or film, but rather a constructed representation of past experience that is influenced by a variety of factors related to the originally encoded event, including general knowledge, personal biases, information from other events, and inferences (Bartlett, 1932; Johnson, 1997; Johnson, Hashtroudi, & Lindsay, 1993; Loftus, 1979, 2003; Roediger, 1996; Schacter, Norman, & Koutstaal, 1998). These constructive processes presumably lead to functionally beneficial representations of the past, but they also cause memory to be prone to error (Schacter, 1999, 2001) The mistaken recognition of an item that is similar, but not identical, to a previously encountered item is a ubiquitous and robust memory error referred to as gist-based false recognition (Koutstaal & Schacter, 1997; cf., Reyna & Brainerd, 1995). For instance, people tend to mistakenly recognize a word that is a synonym of a studied word (Anisfeld & Knapp, 1968), an abstract shape that is structurally similar to studied shapes (Koutstaal, Schacter, Verfaellie, Brenner, & Jackson, 1999; Slotnick & Schacter, 2004), or a pictured object that has the same verbal label as a studied item (Koutstaal, 2006). In these cases, people fail to remember the specific details of an event but can remember more abstracted information — the "gist" — such as the superordinate category of an encountered object.

One possibility is that these errors occur because the original details have been lost from memory, either because they were not encoded originally or because the memory trace has degraded over time. Then, at retrieval, the system relies on more abstract information to reconstruct the lost details. This line of thinking has been implicit in much of the literature. For instance, in the Constructive Memory Framework put

forward by Schacter et al. (1998), it was proposed that gist-based false recognition results primarily from a failure of pattern separation, a process that occurs during encoding. Thus, according to this proposal, gist-based false recognition occurs in large part because the details were not adequately encoded in the first place. Brainerd and Reyna's (2002) Fuzzy Trace Theory proposed that verbatim details are forgotten more rapidly than gist information, a combination that contributes to gist-based false recognition (see also Reyna & Brainerd, 1995).

However, it is also known that people can store an impressive amount of information, particularly about recently encountered pictures or objects. For example, Standing (1973) demonstrated that people could recognize thousands of experimentally presented pictures. More recently, Brady et al. (2008) had participants study 2,896 pictures of objects shown for 3 s each. In a forced-choice test, participants were asked to make subtle distinctions based on memory: the foil was the same object as the studied item, but in a slightly different state (e.g., a bread box with the loaf of bread inside the box or outside the box). Participants scored 87% correct in this condition, suggesting that people can store a large amount of detailed information about recently encountered objects (see also Konkle, Brady, Alvarez, & Oliva, 2010).

These findings present a puzzle. On the one hand, a large body of data on false recognition suggests that the detail stored in episodic memory is limited and that memory relies heavily on constructive processes to compensate for this limitation. On the other hand, the findings concerning highly specific recognition of visual objects suggest that the level of detail stored in episodic memory is far greater than what would have been

expected on the basis of the false recognition findings. Here we attempt to reconcile these seemingly contradictory results.

We suggest that, even when rates of false recognition are high, people do retain many details, but do not adequately utilize them at retrieval. In two experiments, we explored the hypothesis that high rates of false recognition occur when people do not attend to or do not retrieve the relevant perceptual details. To this end, we developed an experimental paradigm that encouraged participants to focus on the relevant perceptual details.

EXPERIMENT 1

Method

Rationale and Design

The conditions of the memory test are depicted in Figure 1. On each trial, the participant was presented with three pictures. Two of the pictures were related to one another because they were both exemplars of the same category and shared a common verbal label. It is important to note that the conditions did not differ systematically in terms of their perceptual presentation; they differed only in terms of the content of the participant's memory. The participant's task was to select one of the items as studied or reject all three items as new ("all new"). In the *baseline target* condition, one of the pictures was a target (studied item) and the other two items were not systematically related to any of the studied items. In the *baseline lure* condition, all three items were not systematically related to any of the studied items and the correct response was "all new". In the *single related item* condition, one of the pictures was related to a studied item. The other two items were not systematically related to any of the studied items. The correct

response was "all new", but we anticipated that participants would falsely recognize the related item with high frequency, reflecting a standard gist-based false recognition effect. In the *target and related item* condition, the target was presented adjacent to the related item. The third item was not systematically related to any of the studied items. In this condition, the nature of the discrimination required was made explicit to the participant. Both the target and the related lure were likely to seem familiar, thus requiring that the participant systematically compared the target and related lure and identified features that distinguished them. If gist-based false recognition occurs because people fail to attend to or retrieve relevant perceptual details still stored in memory, then false recognition rates should be substantially reduced in this condition.

Participants

32 college students (15 male, ages 18-29, mean 22) from the Boston metropolitan area served as participants and were paid $70 (participants were scanned with functional MRI during the experiment; the imaging data will be presented in a separate report). Candidates were excluded for participation that did not meet standard MRI safety criteria, required glasses to see normally, had strabismus or a history of eye surgery, or that were left handed. All participants provided informed consent as approved by the Institutional Review Board at Harvard University. Nine participants were replaced: 4 for poor performance (hits minus false alarms less than .30); 3 for eye tracking problems; 2 for anatomical abnormalities.

Stimulus Materials

384 pairs of object photographs or detailed, colored drawings served as stimuli (Koutstaal, 2006; Koutstaal, et al., 2001). The items within a pair were related to each

other because they were both exemplars of the same category and shared a common verbal label (e.g., wrench, dog, tree). However, the two pictures were perceptually distinct exemplars of the category and, at a minimum, differed in terms of color or orientation. Stimuli were fully counterbalanced (Supplementary Methods).

Procedure

Participants were told that their memory would be tested later and were presented with 144 objects (500 ms duration, 1500 ms ISI) and indicated whether the pictured object could fit into a 13 inch box in the real world by a button press. A box measuring approximately 13 inches was presented. The participant was given a self-paced break halfway through the study session. Then the participant was placed in an MRI scanner. The occurrence of similar foils was clearly explained to the participant. The test was divided into four blocks; each began with 15 s of fixation and ended with 10 s of fixation and contained 12 trials of each condition. Each trial lasted 5 s. With the constraint that the two related items were next to each other, there were four possible arrangements of the pictures; each occurred equally often within each block. 24 fixation trials, also lasting 5 s, were randomly intermixed within each block.

<div align="center">Results</div>

Accuracy

The accuracy data are shown in Figure 2 (reaction times are reported in Supplementary Table 1). Within the baseline target condition, the hit rate was reasonably high (.76, SEM = .02, green bar). Within the baseline foil condition, the correct rejection rate was reasonably high (.69, SEM = .03, blue bar). Therefore, participants performed the task well.

In the single related item condition, participants incorrectly selected the related item at a high rate (.41, SEM = .02, yellow bar). This false alarm rate was approximately four times larger than the baseline false alarm rate to items not presented in a pair, such as the apple in Figure 1 [.11, SEM = .01, orange bar, $t(31) = 18.74$, $p < .001$]. As we expected, there was a robust gist-based false recognition effect in the single related item condition.

In the target and related item condition, the false alarm rate to related foils dropped substantially (.12, SEM = .01, yellow bar), corresponding to a 71% reduction in the rate of gist-based false recognition [$t(31) = 14.97$, $p < .001$]. Nonetheless, within the target and related item condition, false alarm rates to related foils were higher than false alarm rates to unrelated foils [.04, SEM = .01, orange bar, $t(31) = 7.77$, $p < .001$]. Thus, gist-based false recognition was dramatically reduced but not entirely eliminated.

Eye Tracking

To verify that participants systematically compared the target and related foil in the target and related foil condition, we calculated the number of saccades between similar pictures per trial. These data are presented in Figure 3. There were more saccades between related pictures in the target and related item condition (2.80, SEM = 0.13, green bar) than the baseline target condition [1.53, SEM = 0.09, green bar, $t(31) = 18.79$, $p < .001$].

Discussion

Under the same encoding conditions that produced high rates of gist-based false recognition in the single related item condition, participants overwhelmingly chose the correct target in favor of its related foil when given the option to do so. This finding

indicates that the details distinguishing the target from its related foil were, to a large extent, still retained in memory. Hence, gist-based false recognition is not produced by a loss of detail per se. Rather, gist-based false recognition is attributable, at least in part, to a failure to retrieve stored details.

One explanation for our data is that the retrieval failure contributing to gist-based false recognition resulted primarily from a suboptimal deployment of attention: in the single related foil condition, the participant failed to attend to the relevant details; in the target and related foil condition, the participant's attention was drawn to the features that are relevant to the discrimination. However, another critical property distinguishes these conditions: in the target and related item condition, the originally encoded picture is presented at retrieval. In line with the encoding specificity principle (Tulving & Thomson, 1973), we would expect that this reinstatement of the originally encoded information led to increased access to the memory trace. There are thus two distinct mechanisms that could have contributed to the reduction in false recognition observed in Experiment 1: (*i*) attention to detail and (*ii*) reinstatement of detail. In Experiment 2, we attempted to tease apart these two mechanisms.

EXPERIMENT 2

Method

In order to distinguish between the effects of attention and reinstatement, we repeated Experiment 1 with the inclusion of a critical new condition. In the *two related items condition* (Figure 4), the participant was presented with two items, both of which were related to the same study item. The third item in the array was not systematically related to any of the studied items. Both of the related items should seem familiar and the

participant should therefore systematically compare the two related items. Thus, in this condition the participant's attention is drawn to the features that, on average, distinguish exemplars of the category. Critically, however, the originally encoded photograph is not presented.

The methods were identical to those of Experiment 1 except as noted. 30 participants (17 male; ages 18-27, mean 20) were run. Data from 10 participants were replaced: 2 due to movement and 1 due to an anatomical abnormality; 2 for low performance (hit minus false alarm rate less than .30); 4 because they did not produce sufficient gist-based false alarms (at least .25) to enable MRI analysis (note that their inclusion would not alter our conclusions); 1 because he had an unusually low hit rate in the target and related items condition (.30). A new set of 400 object triplets was generated using high quality colored photographs. Participants studied 160 objects, and the memory test was divided into five blocks. Eight trials of each condition occurred in each block. 40 fixation trials occurred in each block.

A parallel experiment was run outside of the MRI environment in order to collect confidence ratings (Supplementary Materials).

<div align="center">Results</div>

The results of Experiment 2 closely replicated Experiment 1. Below, we focus on the novel findings of Experiment 2.

Accuracy

The accuracy data are presented in Figure 5 (reaction time data are presented in Supplementary Table 2). The critical question concerns the gist-based false recognition rate in the two related items condition, where participants attended to the relevant details

but did not benefit from reinstatement of the originally encoded photograph. There was no evidence of a reduction in false recognition to related foils in this condition. In fact, the false alarm rate to related foils was *larger* in the two related items condition (.47, SEM = .02, yellow bar) than in the single related item condition [.38, SEM = .02, yellow bar, $t(29) = 3.51$, $p = .001$]. This result could reflect differences in baseline false alarm rates to foils occurring in pairs and foils not occurring in pairs. Even within the baseline foil condition, the false alarm rate to foils occurring in pairs was larger (.13, SEM = .02, blue bar) than to foils not occurring in pairs [.08, SEM = .01, orange bar, $t(29) = 2.88$, $p <$ .01]. When the false alarm rates to related items were corrected by subtracting the appropriate baseline false alarm rates, there was not a significant difference in false alarm rates to related items in the two related items condition (.33, SEM = .02) and the single related item condition [.30, SEM = .02, $t(29) = 1.38$, $p = .18$].

Eye Tracking

The average number of saccades between related pictures is shown in Figure 6. This measure was larger in the two related items condition (2.98, SEM = 0.11, yellow bar) than in the single related item condition [1.56, SEM = 0.09, yellow bar, $t(29) =$ 13.76, $p < .001$]. These data confirm that participants were directly comparing the two related items in the two related items condition.

<div align="center">Discussion</div>

Experiment 2 differentiated between two mechanisms that could have contributed to the reduction in rates of gist-based false recognition observed in Experiment 1: (*i*) attention to detail and (*ii*) reinstatement of detail. In the two related items condition, participants actively attended to the details that, on average, distinguished between

exemplars of the category. Nonetheless, there was no evidence of a reduction in gist-based false recognition rates. In contrast, in the target and related item condition — where participants had the added benefit of having the originally encoded photograph presented again at retrieval — there was a dramatic reduction in rates of gist-based false recognition, replicating Experiment 1. These results suggest that the outcome of Experiment 1 depended primarily on reinstatement of the studied picture and that increased attention to detail is not sufficient to reduce rates of gist-based false recognition.

One caveat is that the target and related item condition and the two related items condition may have differed in how effectively they drew the participant's attention to specific diagnostic features. For instance, if the target is a red car and both related exemplars are blue cars, then the exemplars will not differ in terms of color in the two related items condition and the participant's attention will not be directed towards color. Nonetheless, the two related items condition does draw the participant's attention to the features that, on average, distinguish exemplars of a category. Therefore, if attention to detail were sufficient to reduce gist-based false recognition, then we would expect some reduction in gist-based false recognition, but we saw no evidence of such a reduction.

It is critical to note that the aforementioned caveat has no bearing on the broader and more significant conclusion that gist-based false recognition results from retrieval failure.

GENERAL DISCUSSION

We investigated the hypothesis that gist-based false recognition does not result from a loss of detail from memory per se, but rather from inadequate utilization of stored

detail during retrieval. Experiment 1 investigated the possibility that people do not attend to the relevant details. In the critical condition of Experiment 1, there was a dramatic reduction in gist-based false recognition. Under the same encoding conditions that produced large rates of gist-based false recognition in the single related item condition, participants overwhelmingly chose the correct target rather than the related foil when given the option to do so. This finding indicates that the detail distinguishing targets and related foils was not lost and could be retrieved under appropriate conditions. Experiment 2 investigated the role of two distinct mechanisms that could have contributed to the initial results: (*i*) attention to detail and (*ii*) reinstatement of detail. The results of the second experiment suggest that the outcome of the first experiment was driven by reinstatement, not attention. These results suggest that retrieval failure is a major factor contributing to gist-based false recognition. The details are, to a large extent, still stored in memory. A critical factor determining whether false recognition will occur is whether the details can be accessed during retrieval.

Our findings are consistent with laboratory experiments on eyewitness identification of suspects in lineups. When the culprit is not present in the line up, rates of mistaken identification are typically high. However, when the culprit is present, participants often choose the culprit correctly (Wells & Olson, 2003). Although research on mistaken eyewitness identification has generally proceeded independently of research on gist-based false recognition, the present results raise the possibility that there are strong parallels between these two types of memory errors and that retrieval failure, in particular, is a common mechanism contributing to both phenomena.

These results offer an interesting contrast to findings obtained using the well-known Deese-Roediger McDerrmott (DRM) paradigm (Deese, 1959; Roediger & Mcdermott, 1995). In this paradigm, participants study a series of associatively related words that are all forward associates of a critical word that was not studied. The standard finding is that people falsely recognize the critical non-studied word very frequently. Schacter, Israel, & Racine (1999) modified the paradigm by presenting a picture with each studied word that depicted what the word named (see also Israel & Schacter, 1997). Participants that studied the pictures showed substantially lowered false recognition rates, relative to participants that didn't study the pictures. Schacter et al. (1999) concluded that participants in the picture condition expected to recollect distinctive pictorial content and demanded access to that content in order to classify a word as "old", thereby allowing them to reject related lures words that did not elicit the expected distinctive information; Schacter et al. (1999) termed this process the *distinctiveness heuristic* (see also Dodson & Schacter, 2002; Gallo, Weiss, & Schacter, 2004). In contrast to the present findings, work on the distinctiveness heuristic suggests that under some circumstances, people can strategically attend to perceptual information to reduce gist-based false recognition**.** A critical difference is that, when using the distinctiveness heuristic, retrieval of *any* pictorial information is sufficient to classify an item as "old". The distinctiveness heuristic requires that people retrieve some sort of distinctive content, but it does not require that the content is highly detailed. In the current experiment, participants must retrieve enough detail to distinguish a target and a related foil. The difficulty of reconstructing this specific visual detail in the absence of the target object contributes to the robustness of gist-based false recognition.

More generally, our findings are consistent with the widely acknowledged fact that false memories are affected by the circumstances at retrieval. With encoding and retention held constant, false memories can be reduced by factors at retrieval, particularly manipulations that encourage participants to monitor contextual details, as has been emphasized in discussion of the source monitoring framework (Dodson & Schacter, 2002; Johnson, 1997; Johnson, et al., 1993; Mather, Henkel, & Johnson, 1997). The current results suggest that, in addition to these previously demonstrated retrieval factors, the failure to retrieve detailed information that is still available in memory is a major contributor to gist-based false recognition.

Our conclusions also have much in common with Fuzzy Trace Theory (Brainerd & Reyna, 2002; Reyna & Brainerd, 1995). According to that model, people encode two independent types of traces: *gist traces* are semantic abstractions and tend to produce false recognition (as well as true recognition); *verbatim traces* contain specific details and support true recognition. It is proposed that verbatim traces are forgotten more rapidly than gist traces, thereby producing such phenomena as gist-based false recognition. However, the model also allows for cue-based retrieval effects. It is postulated that presentation of a target tends to elicit retrieval of a verbatim trace whereas presentation of a related lure tends to elicit retrieval of a gist trace. This notion of differential access to detail depending on the cue information presented to the participant is consistent with our interpretation of the present results. However, rather than focusing on the distinct properties of verbatim and gist traces, we emphasize the role of retrieval cues in providing access to information that is available in memory (Tulving & Pearlstone, 1966;

Tulving & Thompson, 1973) and that can be used under appropriate conditions to avoid gist-based false recognition.

Although these results clearly suggest that retrieval failure is a major factor in gist-based false recognition, it is probably not the only factor. Gist-based false recognition was dramatically reduced in the target and related items condition, but it was not completely eliminated. Even these residual gist-based false memories could be attributable, in part, to retrieval failure: it is in principle possible that the provision of still more cue information, such as reinstatement of the precise study context (Smith & Vela, 2001) or the temporal ordering of studied items (Jacoby, 1972), could reduce the rate of gist-based false alarms further. Nonetheless, it is plausible that the impressive degree of detail stored in memory is not without bound. Our contention is not that loss of detail from memory plays no role whatsoever, but rather that it plays a smaller role than has been previously assumed. In contrast, retrieval failure plays a larger role in generating gist-based false recognition than has been acknowledged to date.

Our results raise the prospect that memory distortions may be avoided by improving the accessibility of information that is still stored in memory. Thus, a pressing avenue for future research is to determine the extent to which similar retrieval failure mechanisms apply to other laboratory demonstrations of false memory, such as imagination inflation (Garry, Manning, Loftus, & Sherman, 1996; Goff & Roediger, 1998), DRM false recognition (Gallo, 2006; Miller & Wolford, 1999; Roediger & Mcdermott, 1995; Weinstein, McDermott, & Chan, 2010), and memory conjunction errors (Reinitz, Lammers, & Cochran, 1992). Further investigation of the role of retrieval failure in producing false memories is likely to help us understand more fully how a

system capable of impressive levels of retention, and which in daily life often serves us well, is also capable of error and distortion (Brainerd & Reyna, 2002; Reyna & Brainerd, 1995; Schacter, 1999).

REFERENCES

Anisfeld, M., & Knapp, M. (1968). Association, synonymity, and directionality in false recognition. *J Exp Psychol, 77*, 171-179.

Bartlett, F. C. (1932). *Remembering; a study in experimental and social psychology*. Cambridge [Eng.]: The University press.

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 of the United States of America, 105*, 14325-14329.

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

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *J Exp Psychol, 58*, 17-22.

Dodson, C. S., & Schacter, D. L. (2002). When false recognition meets metacognition: The distinctiveness heuristic. *Journal of Memory and Language, 46*, 782-803.

Gallo, D. A. (2006). *Associative illusions of memory : false memory research in DRM and related tasks*. New York: Psychology Press.

Gallo, D. A., Weiss, J. A., & Schacter, D. L. (2004). Reducing false recognition with criterial recollection tests: Distinctiveness heuristic versus criterion shifts. *Journal of Memory and Language, 51*, 473-493.

Garry, M., Manning, C. G., Loftus, E. F., & Sherman, S. J. (1996). Imagination inflation: Imagining a childhood event inflates confidence that it occurred. *Psychonomic Bulletin & Review, 3*, 208-214.

Goff, L. M., & Roediger, H. L. (1998). Imagination inflation for action events: Repeated imaginings lead to illusory recollections. *Memory & Cognition, 26*, 20-33.

Israel, L., & Schacter, D. L. (1997). Pictorial encoding reduces false recognition of semantic associates. *Psychonomic Bulletin & Review, 4*, 577-581.

Jacoby, L. L. (1972). Effects of organization on recognition memory. *Journal of Experimental Psychology, 92*, 325-&.

Johnson, M. K. (1997). Source monitoring and memory distortion. *Philosophical Transactions of the Royal Society B-Biological Sciences, 352*, 1733-1745.

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

Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010). Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *Journal of Experimental Psychology-General, 139*, 558-578.

Koutstaal, W. (2006). Flexible remembering. *Psychon Bull Rev, 13*, 84-91.

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

Koutstaal, W., Schacter, D. L., Verfaellie, M., Brenner, C., & Jackson, E. M. (1999). Perceptually based false recognition of novel objects in amnesia: Effects of category size and similarity to category prototypes. *Cognitive Neuropsychology, 16*, 317-341.

Koutstaal, W., Wagner, A. D., Rotte, M., Maril, A., Buckner, R. L., & Schacter, D. L. (2001). Perceptual specificity in visual object priming: functional magnetic

resonance imaging evidence for a laterality difference in fusiform cortex. *Neuropsychologia, 39*, 184-199.

Loftus, E. F. (1979). Malleability of human-memory. *American Scientist, 67*, 312-320.

Loftus, E. F. (2003). Make-believe memories. *American Psychologist, 58*, 867-873.

Mather, M., Henkel, L. A., & Johnson, M. K. (1997). Evaluating characteristics of false memories: remember/know judgments and memory characteristics questionnaire compared. *Mem Cognit, 25*, 826-837.

Miller, M. B., & Wolford, G. L. (1999). Theoretical commentary: The role of criterion shift in false memory. *Psychological Review, 106*, 398-405.

Reinitz, M. T., Lammers, W. J., & Cochran, B. P. (1992). Memory-conjunction errors: miscombination of stored stimulus features can produce illusions of memory. *Mem Cognit, 20*, 1-11.

Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: An interim synthesis. *Learning and Individual Differences, 7*, 1-75.

Roediger, H. L. (1996). Memory illusions. *Journal of Memory and Language, 35*, 76-100.

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*, 803-814.

Schacter, D. L. (1999). The seven sins of memory: Insights from psychology and cognitive neuroscience. *American Psychologist, 54*, 182-203.

Schacter, D. L. (2001). *The seven sins of memory : how the mind forgets and remembers*. Boston: Houghton Mifflin.

Schacter, D. L., Israel, L., & Racine, C. (1999). Suppressing false recognition in younger
and older adults: The distinctiveness heuristic. *Journal of Memory and Language,
40*, 1-24.

Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of
constructive memory. *Annu Rev Psychol, 49*, 289-318.

Slotnick, S. D., & Schacter, D. L. (2004). A sensory signature that distinguishes true from
false memories. *Nat Neurosci, 7*, 664-672.

Smith, S. M., & Vela, E. (2001). Environmental context-dependent memory: A review
and meta-analysis. *Psychonomic Bulletin & Review, 8*, 203-220.

Stampe, D. M. (1993). Heuristic filtering and reliable calibration methods for video-
based pupil-tracking systems. *Behavior Research Methods Instruments &
Computers, 25*, 137-142.

Standing, L. (1973). Learning 10,000 Pictures. *Quarterly Journal of Experimental
Psychology, 25*, 207-222.

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in
episodic memory. *Psychological Review, 80*, 352-373.

Weinstein, Y., McDermott, K. B., & Chan, J. C. K. (2010). True and false memories in
the DRM paradigm on a forced choice test. *Memory, 18*, 375-384.

Wells, G. L., & Olson, E. A. (2003). Eyewitness testimony. *Annual Review of
Psychology, 54*, 277-295.

ACKNOWLEDGEMENTS

FIGURE CAPTIONS

Figure 1. Recognition test conditions in Experiment 1. The participant's task was to select one of the three items as old (studied) or reject all three items as new. The correct response is to select the silver anchor when it is present and reject all three items when it is not. See Rationale and Design.

Figure 2. Accuracy data from Experiment 1. The color of each bar indicates the participant's choice. In the examples provided in Figure 1, the scissors, bulldozers, and beavers are examples of paired unrelated foils. The apple and diskette are examples of single unrelated foils. Error bars show SEM.

Figure 3. Eye tracking data from Experiment 1. The color of each bar indicates the participant's choice. We exclude incorrect responses (with the exception of false alarms to related foils in the single related item condition) because these occurred infrequently and are associated with high estimation error and missing data values for certain participants. Error bars show SEM.

Figure 4. Recognition test conditions in Experiment 2. The task is the same as Experiment 1 (cf. Figure 1), except for the inclusion of the two related items condition.

Figure 5. Accuracy data from Experiment 2. The color of each bar indicates the participant's choice. In the examples provided in Figure 4, the piggybanks, cats, and pretzels are examples of paired unrelated foils. The basketball, accordion, and cow are examples of single unrelated foils. Error bars show SEM.

Figure 6. Eye tracking data from Experiment 2. The color of each bar indicates the

participant's choice. We exclude incorrect responses (with the exception of false alarms

to related foils in the single related item condition and the two related items condition)

because these occurred infrequently and are associated with high estimation error and
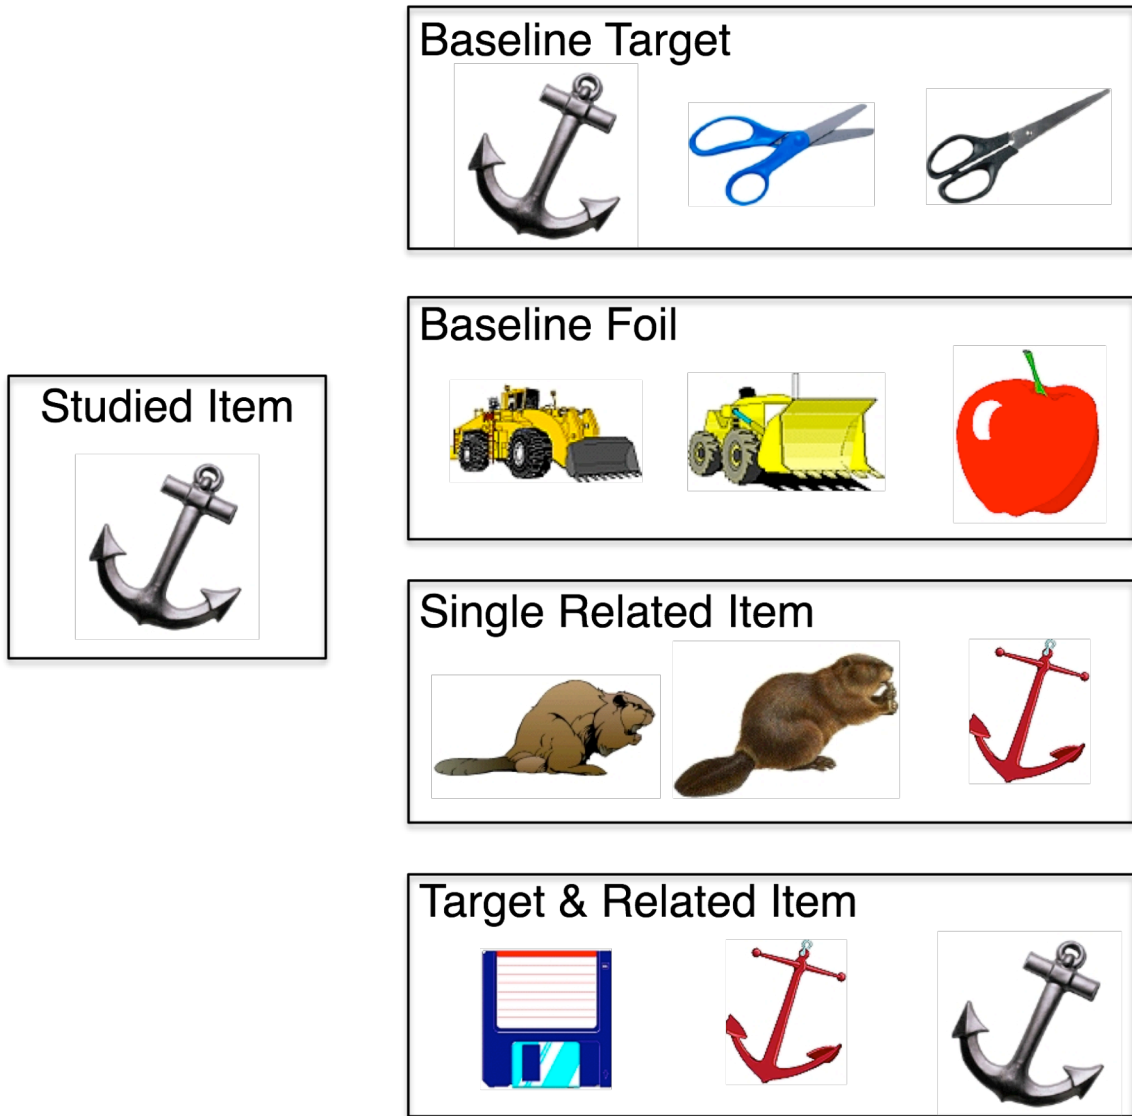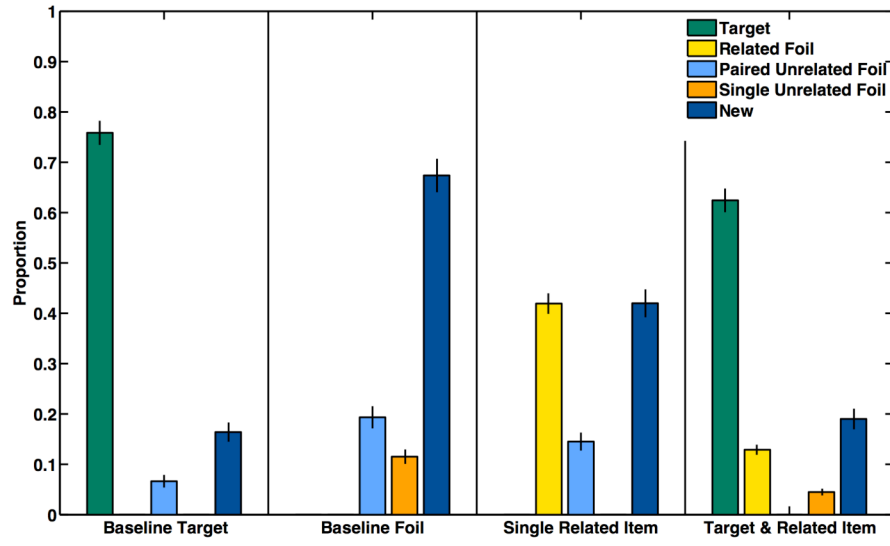
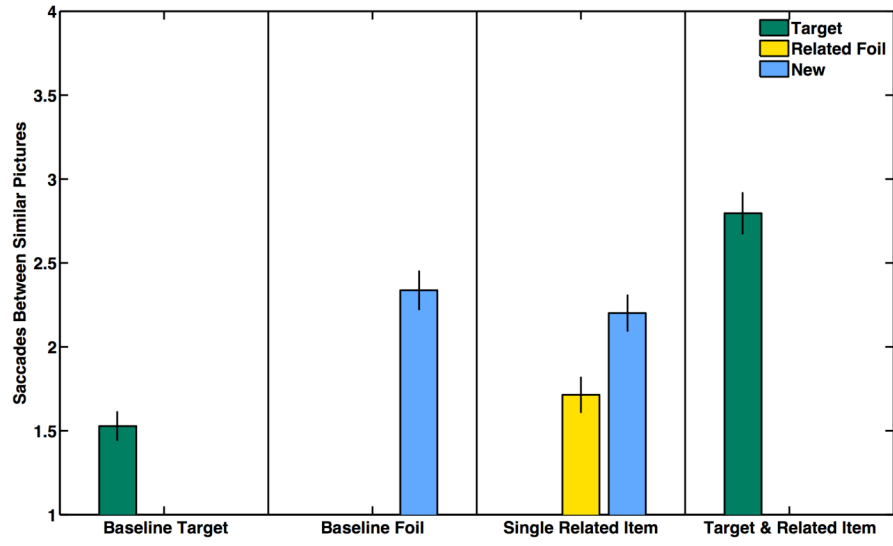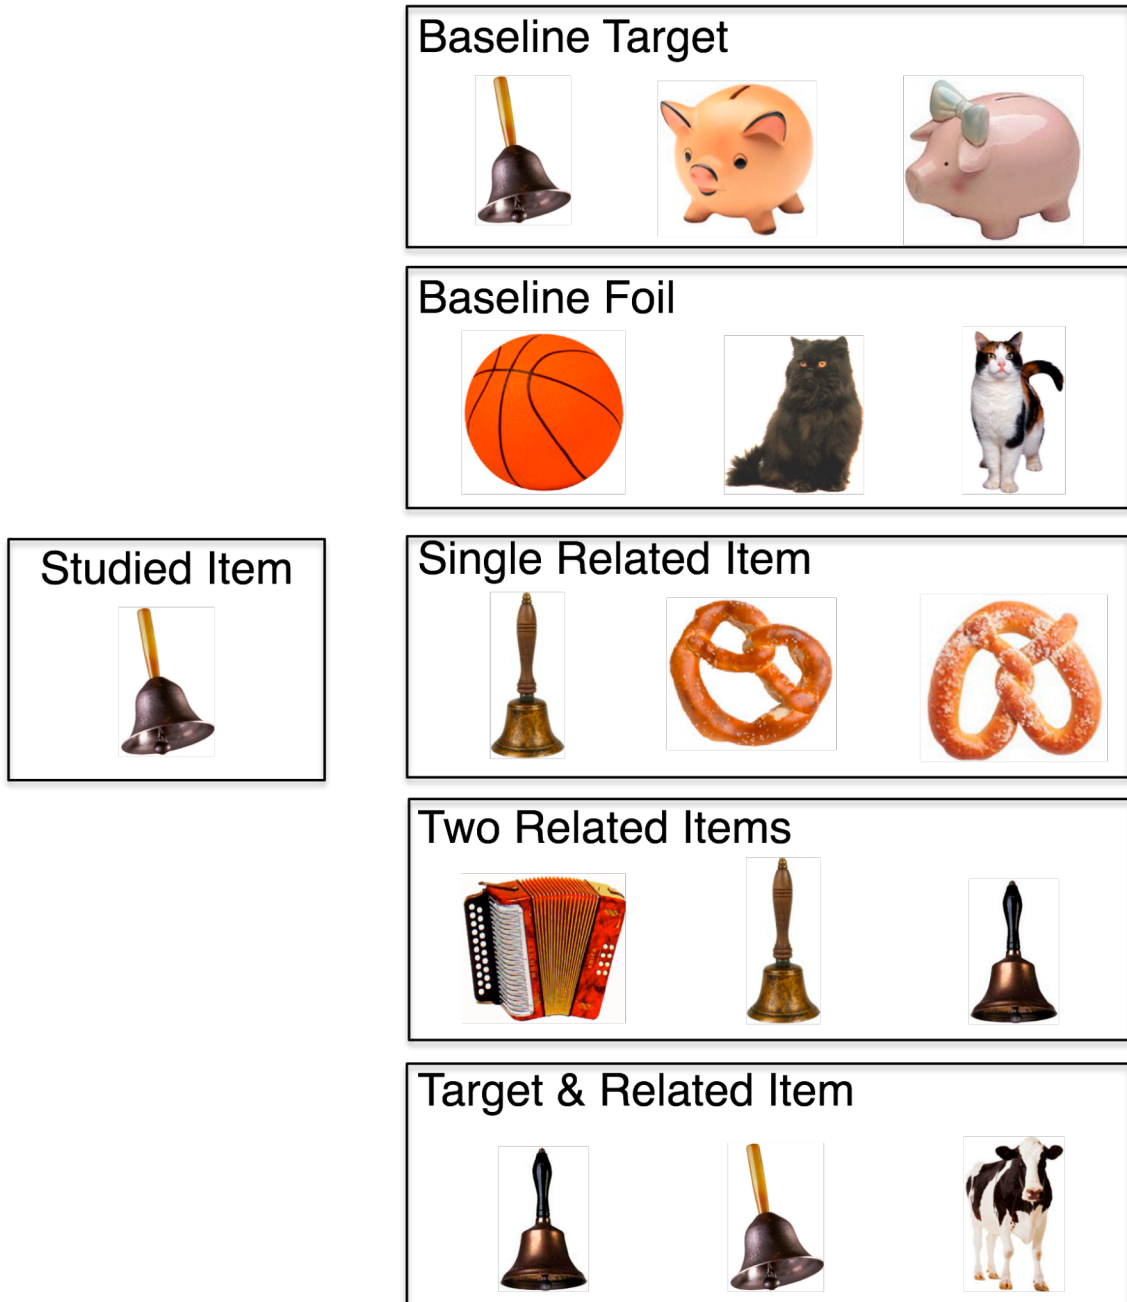missing data values for certain participants. Error bars show SEM.

FIGURE 1

FIGURE 2

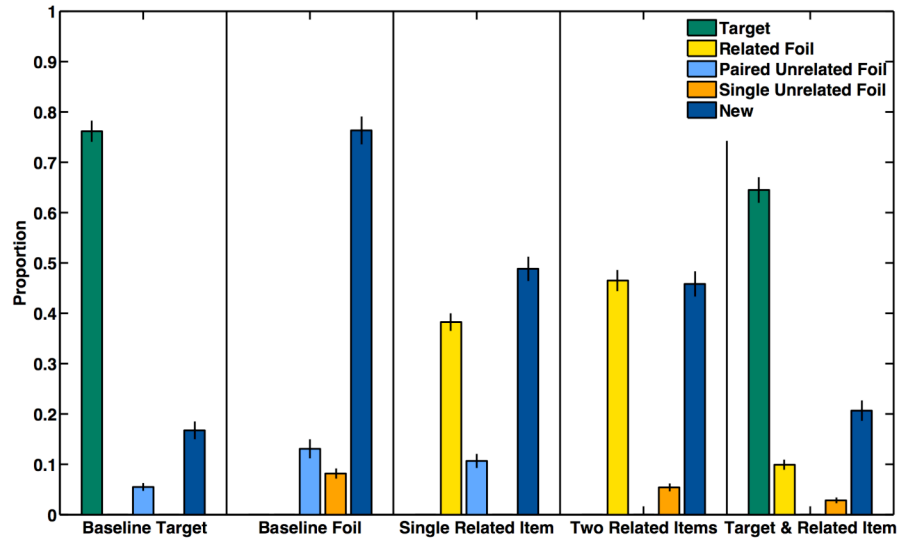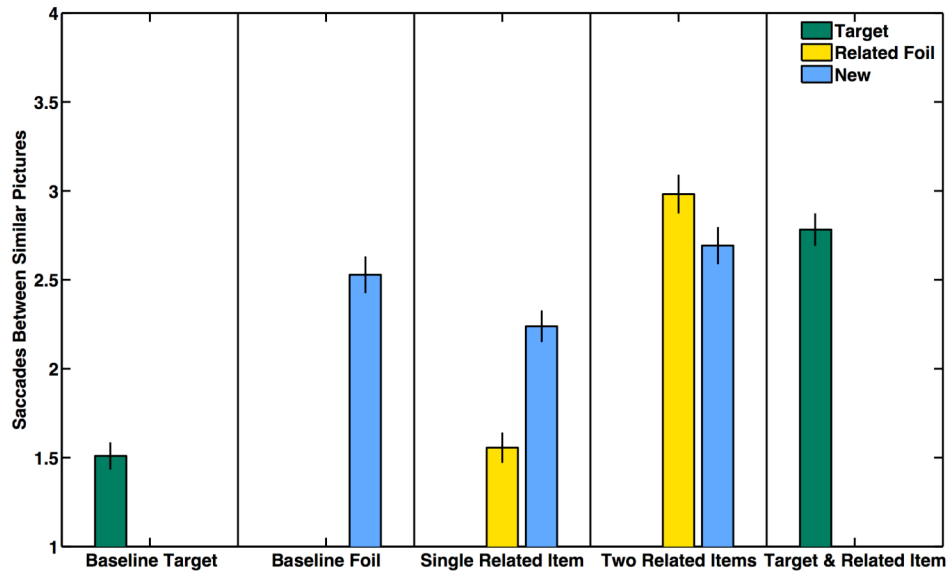FIGURE 3

FIGURE 4

FIGURE 5

FIGURE 6

ONLINE SUPPLEMENTARY MATERIALS

SUPPLEMENTARY METHODS

Stimulus Counterbalancing

Experiment 1. On each trial of the recognition test, two pairs were used to generate the array. One of the pairs was used to produce the two related items and one of the pairs was used to produce the third item (in which case only one of the exemplars in the pair was presented). Thus, with four conditions, there were eight possible assignments of any particular pair. The object pairs were counterbalanced across all eight possible assignments. Then, exemplars serving as targets were further counterbalanced, for a total of 16 iterations of the counterbalancing. This process was repeated to produce a sample size of 32.

Experiment 2. The counterbalancing method used in Experiment 1 was extended to five conditions and three exemplars. With five conditions there were 10 possible assignments of each triplet. Additionally, counterbalancing across three exemplars led to 30 iterations of the counterbalancing.

Eye Tracking

Eye tracking data were collected using an EyeLink 1000 MRI compatible eye tracking system (SR Research Ltd, Mississauga, Ontario, Canada). The pupil and corneal reflection were identified using automated thresholding. The pupil was modeled as an ellipse. For calibration, 9 fixation targets in an equally spaced grid covering the entire stimulus display were presented in a random order. The accuracy of the calibration was validated by presenting an additional 8 fixation targets at new locations on the display as well as the center of the display. The average measured error across the validation points

was typically less than .5° for each participant. Validation was conducted before each block; calibration was rerun as necessary. Monocular data were collected at 1000 Hz. Saccades and fixations were detected online by the EyeLink software using the default thresholds (Stampe, 1993).

SUPPLEMENTARY RESULTS

Experiment 1 Reaction Time.

Reaction time data from Experiment 1 are presented in Supplementary Table 1. Reaction times were obtained by calculating the median reaction time for each participant and then averaging across participants. Two findings are noteworthy. First, gist-based false recognition in the single related item condition was associated with longer reaction times (2031 ms, SEM = 45) than veridical recognition in the baseline target condition [1833 ms, SEM = 34, $t(31) = 5.63$, $p < .001$]. Second, reaction times associated with hits were longer in the target and related item condition (2180 ms, SEM = 76) than in the baseline target condition [1833 ms, SEM = 34, $t(31) = 5.72$ $p < .001$].

Experiment 2 Reaction Time

Reaction time data from Experiment 2 are reported in Supplementary Table 2. Reaction times associated with false alarms to related items in the two-related items condition were longer (2680 ms, SEM = 86) than in the single related item condition [2313 ms, SEM = 74, $t(29) = 6.30$, $p < .001$].

Confidence Experiment

Methods. The procedure was identical to the main experiment except that a confidence scale appeared following each recognition response. Participants responded without a deadline on a four point scale, with 1 labeled "very uncertain", 2 labeled

"somewhat uncertain", 3 labeled "somewhat certain", and 4 labeled "very certain". Eye

tracking data were not collected. 30 participants (9 male; ages 18-27, mean 20) were run.

Data from one participant was replaced due to poor performance (baseline false alarms

greater than baseline hits). Two participants included in the dataset had corrected

recognition (baseline hits minus baseline false alarms) that fell slightly below (.23) the

exclusion criteria for the other experiments (.30). Their exclusion would not alter our

conclusions.

Results. The accuracy data (Supplementary Table 3) and reaction time data

(Supplementary Table 4). from the confidence experiment closely replicated those of

the main experiment. The confidence data are reported in Supplementary Figure 1.

Confidence in gist-based false alarms in the single related item condition was less (2.76,

SEM = 0.07) than in veridical memories during the baseline target condition [3.45, SEM

= 0.06, $t(29) = 9.47$, $p < .001$]. The effects of the attentional manipulation are more

relevant. There was a small but reliable decrease in confidence associated with hits in the

target and related item condition (3.28, SEM = 0.06, green bar) relative to the baseline

target condition [3.45, SEM = 0.06, green bar, $t(29) = 4.03$, $p < .001$]. Similarly, there

was a small but reliable decrease in confidence associated with false alarms to related

foils in the two related items condition (2.47, SEM = 0.08, yellow bar) relative to the

single related item condition [2.76, SEM = 0.07, yellow bar, $t(29) = 4.80$, $p < .001$].

Although rates of gist-based false recognition were not reduced in the two related items

condition, this condition was associated with a modest reduction in confidence.

Supplementary Table 1. Reaction Times (ms) in Experiment 1.

	Baseline Target	Baseline Foil	Single Related Item	Target & Related Item
Target	1833 (34)	NA	NA	2180 (76)
Related Foil	NA	NA	2031 (45)	Low N
New	Low N	2550 (89)	2703 (103)	Low N

Note. We exclude incorrect responses (with the exception of false alarms to related foils in the single related item condition) because these occurred infrequently and are associated with high estimation error and missing data values for certain participants. SEM in parentheses.

Supplementary Table 2. Reaction Times (ms) in Experiment 2.

	Baseline Target	Baseline Foil	Single Related Item	Two Related Items	Target & Related Item
Target	1963 (65)	NA	NA	NA	2105 (79)
Related Foil	NA	NA	2313 (74)	2680 (86)	Low N
New	Low N	2615 (102)	2761 (108)	2811 (99)	Low N

Note. We exclude incorrect responses (with the exception of false alarms to related foils in the single related item condition and the two related items condition) because these occurred infrequently and are associated with high estimation error and missing data values for certain participants. SEM in parentheses.

Supplementary Table 3. Accuracy in the Confidence Experiment.

	Baseline Target	Baseline Foil	Single Related Item	Two Related Items	Target & Related Item
Target	.73 (.03)	NA	NA	NA	.66 (.03)
Related Foil	NA	NA	.31 (.02)	.37 (.02)	.08 (.01)
Paired Unrelated Foil	.06 (.01)	.13 (.02)	.11 (.02)	NA	NA
Single Unrelated Foil	NA	.08 (.01)	NA	.05 (.01)	.03 (.01)
New	.19 (.02)	.76 (.04)	.55 (.04)	.55 (.03)	.21 (.02)

Supplementary Table 4. Reaction times (ms) in the Confidence Experiment.

	Baseline Target	Baseline Foil	Single Related Item	Two Related Items	Target & Related Item
Target	2169 (50)	NA	NA	NA	2408 (74)
Related Foil	NA	NA	2508 (74)	3003 (86)	Low N
New	Low N	2619 (82)	2769 (83)	3245 (210)	Low N

Supplementary Figure 1. Confidence data from Experiment 2, collected in a separate group of 30 participants run outside the MRI environment. The color of each bar indicates the participant's choice. We exclude incorrect responses (with the exception of false alarms to related foils in the single related item condition and the two related items condition) because these occurred infrequently and are associated with high estimation error and missing data values for certain participants. Error bars show SEM.