



Misremembering

Sarah K. Robins

To cite this article: Sarah K. Robins (2016) Misremembering, *Philosophical Psychology*, 29:3, 432-447, DOI: [10.1080/09515089.2015.1113245](https://doi.org/10.1080/09515089.2015.1113245)

To link to this article: <http://dx.doi.org/10.1080/09515089.2015.1113245>



Published online: 22 Feb 2016.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

Misremembering

Sarah K. Robins

University of Kansas, Department of Philosophy, 3090 Wescoe Hall, 1445 Jayhawk Blvd., Lawrence, KS 66045, USA

ABSTRACT

The Archival and Constructive views of memory offer contrasting characterizations of remembering and its relation to memory errors. I evaluate the descriptive adequacy of each by offering a close analysis of one of the most prominent experimental techniques by which memory errors are elicited—the Deese-Roediger-McDermott (DRM) paradigm (Deese, 1959; Roediger & McDermott, 1995). Explaining the DRM effect requires appreciating it as a distinct form of memory error, which I refer to as misremembering. Misremembering is a memory error that relies on successful retention of the targeted event. It differs from both successful remembering and from confabulation errors, where the representation produced is wholly inaccurate. As I show, neither the Archival nor the Constructive View can account for the DRM effect because they are insensitive to misremembering's unique explanatory demands. Fortunately, the explanatory limitations of the Archival and Constructive Views are complementary. This suggests a way forward. Explaining misremembering—including how it differs from both successful remembering and confabulation—requires a hybrid theory of memory, combining the Archival commitment to discrete retention with the Constructive approach to retrieval. I conclude the paper with the beginning sketches of such an account.

KEYWORDS

Constructive memory; DRM effect; memory; memory errors

1. Introduction

Philosophical theories of the nature of memory can be sorted into two general types: traditional *Archival* views and contemporary *Constructive* ones. They offer contrasting characterizations of remembering and its relation to memory errors. According to the Archival picture, memory is a preservative capacity that stores discrete representations of particular past events. Memory errors are understood, correspondingly, as malfunctions in the process of retrieving these representations. Constructivists argue that evidence from memory science tells against the Archival View. Errors pervade everyday remembering, occurring far too frequently to be considered malfunctions of a preservative process. Constructivists thus offer an alternative account, according to which all attempts at remembering—both successes and errors—are outputs of a single, adaptive process by which plausible representations are constructed at the time of recall (De Brigard, 2014; Michaelian, 2012; Sutton, 2007).

Since it views the Archival view as descriptively inadequate, Constructivism's viability as an alternative depends upon its ability to account for the range of errors revealed by memory science. While the view is both intriguing and promising, it cannot—in its current form—meet this challenge. Thus far, the Constructivist approach has been to address memory errors en masse, appealing to one general

constructive process to capture them all. This leaves the Constructive View insensitive to the unique explanatory demands of particular types of memory error. Accounting for these errors will require a hybrid theory of memory that combines elements of both the Archival and Constructive views, or so I shall argue.

I defend these claims by offering a close analysis of one of the most prominent experimental techniques by which memory errors are elicited—the Deese-Roediger-McDermott (DRM) paradigm (Deese, 1959; Roediger & McDermott, 1995). In DRM experiments, participants falsely “recognize” items as belonging to a previously presented set when the items are similar to those in the set. The DRM effect is one of the best-established memory errors in the empirical literature, providing a crucial test of the descriptive adequacy of both the Archival and Constructive views.

Explaining the DRM effect requires appreciating the distinct form of memory error involved. I refer to such errors as *misrememberings* and propose an account of misremembering as the first step in building a much-needed taxonomy of memory errors. Misremembering is a memory error that relies on successful retention of the targeted event.

With the nature of misremembering clear, I turn to an evaluation of the Archival and Constructive views. As I show, neither can account for the DRM effect. They are insensitive to misremembering’s unique explanatory demands. The Archivalist is unable to explain how one could retain information from a past event and yet fail to produce accurate representations of that event during retrieval. The Constructivist provides a rich account of the influences on the representation constructed at retrieval, but cannot explain the DRM effect because it fails to distinguish between cases of misremembering and confabulation. Fortunately, the explanatory limitations of the Archival and Constructive Views are complementary, suggesting a way forward via a hybrid theory. I conclude the paper with a discussion of ways that such a hybrid view could be developed.

2. Memory Errors: The DRM Paradigm and Misremembering

A theory of remembering must account for memory’s successes and failures. Doing so requires an understanding of the kinds of memory errors that are possible. Section 2.1 provides an extended discussion of one such memory error—the DRM effect. In section 2.2, I argue that the DRM effect should be understood as misremembering, a particular type of memory error.

2.1. The DRM Paradigm

The DRM Paradigm was introduced by Deese (1959) and revived by Roediger and McDermott (1995).¹ Participants are presented with a list of related items to memorize. The most basic form of the task involves a list of semantically related words, such as *nurse*, *sick*, *medicine*, *clinic*, *health*, and *hospital*. After a short break, participants are asked whether they recognize some additional items as members of the initial list. The recognition test involves three types of items:

1. List items (e.g., *hospital*)
2. Unrelated, non-list items (e.g., *apple*), and
3. Related, non-list items (e.g., *doctor*)

When asked whether they recognize *hospital*-type items, participant responses follow a standard serial position curve: the recognition rate varies as a function of the item’s location on the list. Participants are most likely to recognize words at the beginning or end of the list—doing so 70–80% of the time. Words in medial positions are recognized only 40–50% of the time. Participants rarely report recognizing *apple*-type items, doing so less than 10% of the time. In the DRM task, questions of the third type are of most interest. Do participants treat *doctor*-type items more like *hospital* or more like *apple*? Responses are most similar to *hospital*; participants report recognizing *doctor*-type items 70% to 80% of the time, comparable to the best-recognized items from the initial list.

One might be inclined to dismiss this result as the product of a highly contrived and artificial task. But this interpretation is difficult to maintain. First, the DRM effect has been replicated extensively. According to a recent estimate, there has been “one new [DRM] experiment every 2 weeks for the past 10 years” (Gallo, 2006, p. 22). The effect can also be obtained with a range of stimuli. False recognition of *doctor*-type items occurs when the set of words are semantically related, as in the above example, but also when the words are related by category, phonology, or orthography.² So too, the effect persists with recall tests and can be found with non-linguistic items, such as pictures, faces, and dot arrays.³

The DRM effect is both pervasive and persistent.⁴ Standard versions separate learning and recognition by only a few minutes, but the effect remains when they are separated by hours, days—and in some cases—months (see Anastasi, Rhodes, & Burns, 2000). Even more strikingly, confidence in one’s recognition of misremembered items continues despite warnings to be vigilant against such errors, regardless of whether the warning is given during learning or recognition.

2.2. The DRM Effect as Misremembering

The DRM effect is perplexing because it involves both success and failure. In making this error, participants both forget and remember what was on the list. When a participant claims to recognize *doctor*, she makes an error. *Doctor* was not on the list. But to claim that *doctor* and *hospital* were on the list, while *apple* was not, indicates successful retention. Participants are able to make this error only if they retain at least some of the information from the initial list.

Such mistakes are indicative of a distinctive type of memory error, one that I suggest is best termed *misremembering*. In discussions of memory errors, “misremembering” “false memory,” and “confabulation” are often used interchangeably. I recommend a departure from this practice, using each term to capture a particular failure of remembering. I focus on misremembering, defining it as follows:

Misremembering is a memory error that relies on successful retention of the targeted event. When a person misremembers, her report is inaccurate, yet this inaccuracy is explicable only on the assumption that she has retained information from the event her representation mischaracterizes.

Not all memory errors are misrememberings. There are cases of forgetting (temporary and permanent) and confabulation, to name only a few. The distinction between misremembering and confabulation is particularly important. Whereas misrememberings result from the distortion of retained information, confabulations are wholly inaccurate, reflecting no influence of retained information from a particular past event. Suggestibility studies, pioneered by Elizabeth Loftus, are examples of confabulation. These studies show that, as a result of mildly suggestive questioning, participants can come to “remember” events they never experienced, such as being lost in a shopping mall as a small child or having been hospitalized overnight (Loftus & Pickrell, 1995). The confabulated “memories” draw on information from many sources; their explanation does not require appeal to a particular past event that has been distorted, as is required for the DRM effect.

The distinction between misremembering and confabulation runs analogously to that between illusion and hallucination in the study of perception.⁵ In illusory perception, one perceives an object as having properties that it does not. Hallucination is a more extensive error; it occurs when the entire perceptual experience—both *what* is perceived and *how* it is perceived—is illusive. In the DRM paradigm, *doctor*-type items produce a mnemonic illusion. Loftus’ results are more akin to hallucinations. Of course, the distinction between these forms of error need not be categorical. The difference between misremembering and confabulation may, ultimately, be one of degree.⁶

Misremembering errors are not restricted to laboratory contexts. They are familiar from everyday experience. In reflecting on his tenth birthday party, for example, Ted might recall that he received a green bike rather than a blue one or report that Peggy served the punch, when it was her sister, Joan, who did. Swapping green for blue or Peggy for Joan is a mistake, but the error depends on the retained information.

Moreover, the DRM effect is analogous to results found in experimental studies of memories for significant cultural events, such as the fall of the Berlin Wall, the terrorist attacks of September 11th,

and the Arab Spring protests in Egypt. While these memories are often referred to as “flashbulb memories,” in homage to their seeming indelibility, studies reveal that the details presented in these recollections change over time (e.g., Neisser & Harsch, 1992). As in the DRM task, participants are often confident in their reports of details that were not part of their previous experience and, in some cases, retain this confidence even in light of evidence that tells against the accuracy of their retelling.

I turn now to a review of the Archival and Constructive accounts of memory to see whether either has the resources to explain this memory error.

3. The Archival View of Memory

The Archival View is the traditional view of memory within both philosophy and psychology.⁷ Rarely does it receive an explicit and elaborate defense. In what follows, I identify and discuss the three main tenets of the Archival View.

3.1. Memory as a Preservative Capacity

According to the Archival View, the aim of memory is to preserve one’s previous experiences for subsequent remembering. It does so via memory traces: detailed representations of particular past events retained and available for later recall. The Archival View is often traced back to Aristotle and his characterization of memory as the preservation of perception. It is also associated with empiricist theories of mind in the 17th and 18th centuries, especially those of Locke and Hume. Proponents of the Causal Theory of Memory distance themselves from the account of memory traces found in these earlier accounts (e.g., Bernecker, 2010; Debus, 2010; Martin & Deutscher, 1966), but nonetheless remain committed to the Archival framework.

Broader commitment to the Archival View is seen in memory’s recruitment to address other philosophical puzzles, such as those that arise in discussions of knowledge, personal identity, and agency. Memory is appealed to as a faculty through which our experiences leave behind impressions that are stored and later revived. It serves as a psychological constant that connects beliefs, experiences, intentions, obligations, commitments, and the like across time.

Memory scientists, too, have traditionally endorsed a form of the Archival View. The scientific study of memory is organized around the idea that remembering involves three distinct processes: encoding, storage, and retrieval. “Encoding” is the process by which information makes its way into memory. “Storage” is process by which information is maintained in memory, and “retrieval” is the process of recovering that information. Although challenged by the recent trend toward Constructivism, the Encoding-Storage-Retrieval model remains the standard information-processing framework for memory.

3.2. Remembering as Retrieval

The preservationist picture of memory’s function yields a straightforward account of remembering: it is an act of retrieval. If memory creates and retains traces of past events, then remembering is a matter of retrieving the trace corresponding to the desired event. Much as one might fetch a book from the library by jotting down its call number and walking the stacks to reach it, so too remembering is characterized as a search of one’s mental repository. Aristotle, for example, describes recollection as “searching for an image in a corporeal substrate” (*Parva Naturalia*, 453a15–20). Similarly, William James claims that “we search in our memory for a forgotten idea just as we rummage our house for a lost object” (1890, p. 654).

Scientific models share this commitment. As discussed in section 3.1, the third stage of memory is labeled “retrieval” and understood as a process of scanning the memory store.⁸ Each encoded experience results in a discrete representation housed in the memory store. Its location is recorded so that retrieval can later be achieved by consulting this address.

3.3. Memory Errors as Malfunctions

Although proponents of the Archival View focus on explaining cases of successful remembering, they are not committed to memory's infallibility. There is, however, a close connection between how Archivalists view the nature of memory and the kinds of errors considered to be possible. Archivalists view memory errors as malfunctions of a preservative process designed to result in successful retrieval. Remembering requires searching for, and locating, a memory trace corresponding to the event one intends to remember. At each stage there is an opportunity for malfunction. It could occur during encoding or storage: the trace could decay, be misplaced, or never have been formed. The retrieval process could go awry if the wrong location is searched or if the item is not recognized. Errors could also result from failure to access the memory system or from random, inexplicable variance in this otherwise reliable process. The result is four general types of malfunction explored in the next section.

4. The Archival View and the DRM

The Archival View allows for four general types of memory error: an error could result from guessing (section 4.1), procedural error (section 4.2), trace decay (section 4.3), or noise (section 4.4).⁹ In what follows, I consider each of the Archivalist's explanatory strategies and show why the DRM effect resists their capture.

4.1. DRM Effect as Guessing

The DRM effect could be explained as use of general reasoning rather than memory. After all, participants do not have to remember hearing *doctor* in order to answer the recognition question. Their response may reflect a guess that *doctor* was likely to have been on the list.

There are, however, two reasons to reject this interpretation. First, it conflicts with participant self reports. DRM experiments often include *remember/know* judgments. Participants are instructed to respond "remember" if they can recall specific, vivid details about the item and to respond "know" if they are guessing. For *doctor*-type items, participants select 'remember' as often as for *hospital*-type items (58% and 57% of the time, respectively, in Roediger and McDermott's [1995] initial study).¹⁰ In fact, participants often provide details of what they "remember" about the non-presented word, including where it occurred on the list, what they were thinking when they heard it, and what the voice reading the word sounded like. One participant, for instance, claimed that the non-presented word *piano* caused him to form "an image of trying to get a grand piano through the front door at home" (Dewhurst & Farrand, 2004, p. 408).

Second, even if one is reluctant to trust participant reports, patterns in recognition behavior challenge the idea that they are using a general reasoning strategy. If they were simply guessing, then the overall rate of recognition should be indistinguishable from chance. But participants are more likely to recognize items that were on the list (e.g., *hospital*) and less likely to recognize items that were not (e.g., *apple*)—reflecting at least partial reliance on memory. Further, while participant responses are influenced by the overall frequency of the items they are asked to recognize, this alone cannot explain recognition rates. The influence of word frequency is mediated by the strength of the association between the non-presented item and list items (Gallo & Roediger, 2003). That is, participants' recognition reports depend on the items presented during list-learning.

4.2. DRM Effect as Procedural Error

The Archivalist's second option is to characterize the DRM effect as a procedural error, as a malfunction in the search process. The procedure could fail to specify the right search location or could go in search of the wrong item.

The DRM effect cannot be explained as either of these. If retrieval led to the wrong location, one would expect an overall decrease in recognition performance. But participant responses for presented items follow a standard serial position curve reflective of remembering. It is also difficult to see how the DRM effect could be explained as a search for the wrong item, given that participants are often making recognition judgments. In fact, the DRM effect seems to require the absence of procedural error. If the search did not locate the list, then how could list items influence participant responses? DRM errors are not procedural errors.

A defender of the Archival View may prefer an alternative interpretation. Participant reports could reflect successful remembering of an error during the initial encoding. That is, if a participant realized the list contained items related to *doctor*, he or she might represent the list as a set of *doctor*-related items and go on to confuse this elaboration on the list for recognition of an item from the list at the time of recall. This would be an encoding error, but not a memory error, making retrieval of the non-presented item explicable.

Some errors could be explained in this way, but the tactic is of limited use. It is difficult to reconcile with the persistence of the DRM effect after explicit warnings (e.g., McDermott & Roediger, 1998). Further, if recognition of *doctor* is merely an elaboration on the list, then why do participants report vivid details about this word? Studies of false memory in other contexts cast further doubt on this interpretation, by showing that participants falsely recognize or recall past events even when the information that prompts the error is only made available at retrieval.

4.3. DRM Effect as Trace Decay

The DRM effect may yet be explicable if it can be explained as the result of trace decay. Archivalists allow that not all information gets into memory, and even for the information that does, it may become degraded over time.

However, decay is unlikely to produce the DRM effect. Loss of the trace for the list-learning event should result in forgetting, rendering the participant unable to recognize not only *doctor*, but *hospital* as well. But the DRM effect is an error of over-, not under-, reporting on the contents of memory. Decay alone is insufficient to explain false recognition of *doctor*-type items.

Proponents of the Archival View might suggest that the error results from trace decay combine with the influence of general reasoning. If some list items have decayed, the participant might infer that *doctor* was on the list because of its relation to the remaining items. But again, this result does not comport well with the full range of results. As pointed out earlier, participants take themselves to be remembering, not guessing, when they report recognizing *doctor*. The decay-plus-inference interpretation cannot explain why participants create—and believe—stories about the presentation of non-presented words. Further, if instances of false recognition were the result of reasoning's intervention, then one would predict differences in the response times for cases of recognition and inference, but there are no such systematic differences (Gallo, 2006).

4.4. DRM Effect as Noise

The Archivalist has one possible explanation of the DRM effect remaining. Some memory errors may be the result of noise—random, inexplicable variance in the retrieval process. Anderson's ACT-R model of cognitive processing, for example, attempts to explain memory errors in this way (e.g., Anderson, Bothell, LeBièrè, & Matessa, 1998). The DRM effect could be one such error.

This interpretation of the DRM should be unattractive to the Archivalist for reasons enumerated above. The DRM effect is not a random error; it is highly systematic. There is a straightforward explanation of why participants claim to recognize words like *doctor*, but not *apple*. To appeal to inexplicable variance in explaining the DRM effect would require downplaying critical features of the experimental design.

The DRM effect cannot be reconciled with the Archivalist's characterization of memory errors as malfunctions. The inability to accommodate this effect reveals a serious limitation of the Archival approach.

5. The Constructive View of Memory

The Constructive View offers an alternative approach to remembering, designed to account for the preponderance of memory errors, like the DRM effect, in contemporary psychology. The Constructive View has been available in psychology for most of the 20th century (e.g., Bartlett, 1932; Neisser, 1967). It is only more recently—as evidence of memory errors has accumulated—that it has risen to prominence (Klein, 2013; Loftus, 2003; Schacter, Norman, & Koustaal, 1998).

Philosophical accounts of constructive memory, which are my primary focus below, urge a rethinking of how memory achieves its function in response to the nature and quantity of memory errors revealed by memory science (De Brigard, 2014; Michaelian, 2012; Sutton, 2007).¹¹ The worry is not simply that the Archival View leaves such errors unexplained. Rather, it is that the extent of these errors conflicts with the Archival understanding of memory's function. In light of the experimental evidence, the Archivalist is forced to conclude that memory rarely functions properly. A better option, the Constructivist claims, is to think that memory's preservative function has been misunderstood.

Amongst Constructivists, there are differences in how the contrast from the Archival View is drawn. But all versions of Constructivism share two general features. Below I discuss these features and the three most prominent versions of Constructivism—Connectionist, Gist-Based, and Episodic Hypothetical Thinking accounts.

5.1. Generalizing Memory's Architecture

In response to evidence of memory's malleability, Constructivists reject the Archival View's characterization of how memory preserves information from the past. Memory does not appear to store discrete representations of particular past events. Memory science has shown, repeatedly, that the representations produced in remembering contain information from a range of sources. Sometimes memory of a past event reflects the rememberer's general background knowledge and cultural assumptions, as Bartlett (1932) showed in a set of studies often identified as the beginning of Constructivism. More recent studies reveal that people retain only a rough sketch of past events that reports generalized interpretations of the event that favor their future aims and interests (Brainerd & Reyna, 2005). In some cases, what's recalled is not a generalization, but a representation that swaps details of one event for another, or combines information from several distinct events into a single, "episodic" memory (Neisser, 1981).

Attempts at recall are also influenced by the conditions of retrieval. Constructivists note that memories do not simply reappear, as the Archivalist suggests. Instead, the representation produced is influenced by features of the retrieval context, including the cue used to prompt memory, the rememberer's current emotional state, and information available in the environment. The beliefs and values that one holds at the time of remembering, for example, exert an influence on how one recalls past beliefs and values (Ross, 1989).

To be clear, Constructivists do not deny that people retain information about the past. The question is whether explaining how we do so requires positing discrete representations of particular past events. Since memory's outputs are amalgamations of information from several sources, it is reasonable to conclude that the output reflects the organization of the underlying memory store. And so Constructivists recommend replacing the Archivalist repository with a more generalized or distributed form of information storage. Versions of Constructivism can be distinguished by the specific architecture endorsed.

5.1.1. Connectionism

Some Constructivists reject the Archivalist focus on memory as a capacity for preserving *particulars*, suggesting instead that memory preserves *patterns*. This version of Constructivism is often referred to as *Connectionism* (Sutton, 2007), which can be misleading. The form of connectionism at issue here is best understood as an account of semantic networks connected by patterns of spreading activation (e.g., Bechtel & Abrahamsen, 2002; Plaut, 1995). Connectionist principles are applied to *memories*, not mental content. Information acquired as the result of a particular event is distributed throughout the network, as a particular pattern of associations between nodes, but basic concepts are represented as discrete, semantically evaluable nodes.

For Connectionists, memory is characterized as a general network of information, which contains a node corresponding to each basic idea that the subject has and, in turn, these nodes are connected to one another in ways that reflect associations between these ideas. Each node then has an activation level—or strength—reflecting how often and how recently the subject has encountered the corresponding idea. When a particular node is active, this activation spreads through the network via its connections to other nodes.

On this view, memory encodes information from events, but does so in a way that privileges similarities across experiences over the details of any one experience. A person may, for example, attend several birthday parties over the course of her life. As a result of these experiences, she will remember general features of such events—that parties often involve cake and presents, say—rather than remembering the specific cake served or presents received at any particular party.

5.1.2. Gist-Based

Whereas Connectionists characterize memory as generalizing across events, proponents of Gist-based views focus on the generalizations that occur within an individual memory (e.g., Michaelian, 2012).¹² On this view, memory keeps track of particular past events, but the aim is not to provide a pristine and detailed record of each event. Instead, memory provides a rough-and-ready summary—the gist—of the experience. As Michaelian explains the view, there is “no literal record of experience” (2012, p. 329). The process is constructive because the memory of each event undergoes a series of transformations that remove detail and standardize the content. Take, for example, your memory of a recent conversation with a friend. Likely you can recall general features of that event: topics of conversation, your friend’s mood, time of day, and so forth. It is unlikely, however, that you remember many details—the precise words and phrases used, the number of hand gestures made, and so on.

Proponents of Gist-based accounts are quick to emphasize the advantages this form of storage affords the rememberer. Retaining detailed accounts of every past experience would be extraordinarily demanding in terms of information storage and processing. What’s more, these demands would meet with little payoff, since most event details are unlikely to be of future use. The Gist-based view is understood to be more adaptive because it suggests that memory retains only what is likely to be useful for the would-be rememberer’s future thoughts, plans, and actions.

5.1.3. Episodic Hypothetical Reasoning

Other versions of Constructivism press the denial of the Archival View of memory’s function, rejecting the claim that memory is a distinct cognitive capacity (De Brigard, 2014; Schacter & Addis, 2007). It is, instead, only one way of using a more general cognitive process of inference. De Brigard characterizes it as *Episodic Hypothetical Thinking*: a system for creating “self-referential mental simulations about what happened, may happen, and could have happened to oneself” (2014, pp. 174–175). Remembering the past, making decisions in the present, and planning for the future all rely on the same underlying process, operating over a single network of information. They are distinguished only by the tense assigned to the outcome.

Proponents of this view highlight the functional and anatomical similarities between counterfactual reasoning, future planning, and episodic remembering. Each process can be understood as a type of mental simulation, whereby patterns of information are flexibly combined and recombined to produce

the desired representations. Neuroimaging studies reveal that the same underlying core network of brain processes is active during remembering, imagining, planning, and hypothetical thinking (De Brigard et al., 2013), further discrediting appeals to a distinct memory store.

5.2. Remembering is Reconstruction

Since Constructivists deny the existence of discrete representations of particular past events, they also reject the account of remembering as an act of retrieval. Instead, Constructivists view remembering as a process of building a representation of the event one is trying to remember using all available sources of information. It is “an inferential process, constructive not reproductive” (Sutton, 2007, p. 219).¹³ Each version of Constructivism comes with its own account of this inferential process to suit the underlying architecture from which the representations are built.

5.2.1. Connectionism

For Connectionists, the network’s activation level at any given time is a reflection of how often and how recently the would-be rememberer has encountered the ideas corresponding to the nodes in this network. The network state guides the constructive act of remembering. Construction begins from the node(s) associated with the cue or prompt and spreads via previously established connections. The nodes that are most easily accessible from the cue-node will be the ones that have been most recently co-active with the cue (the most recent) and those that are, overall, best connected to the cue (the most frequent). When trying to recall a past event, the information from the activated nodes is used to construct the representation. In recognition tasks, when a person is trying to determine whether a prompt or cue was part of a past experience, a person’s responses are based on the ease with which such information can be accessed (i.e., information that is easily called to mind is interpreted as recognized; information that is more difficult to call to mind is not).

5.2.2. Gist-Based

Gist-based views of remembering claim that remembering relies on a retained gist, or “episodic interpretations of concepts” (Brainerd & Reyna, 2002, p. 165), which summarize the event, becoming more generalized and abstract over time.¹⁴ When a person is asked to recall a past event, he or she uses the schematic information contained in the gist to construct a plausible representation. If, for example, the event is a birthday party, then the representation will include details compatible with one’s schema for birthday parties, such as presents and cake. When a person is asked whether he or she recognizes something from a past event, the rememberer’s task is to determine whether this cue is consistent with the retained gist.

5.2.3. Episodic Hypothetical Thinking

According to De Brigard’s Episodic Hypothetical Thinking account, remembering is but one way of using a general system for hypothetical reasoning. When characterizing this underlying network of information, De Brigard appeals to structural features similar to those found in both Connectionist and Gist-based views. The network is described as an amalgamation of prior knowledge and schematic generalizations. Patterns in the network are determined by the agent’s expertise or relative frequency of exposure to the items represented as nodes in this network (De Brigard, 2014, p. 170). This network is used to generate inferences about the plausibility or likelihood of various scenarios based on the patterns established in the network. In remembering, the network is used to construct a representation of what could have happened or what was most likely to have happened. In recognition, the patterns in the network are used to assess the plausibility or likelihood of the information from the prompt.

5.3. Memory Errors Are Not Malfunctions

Regardless of which inferential principle they favor, all Constructivists stress the similarity between the processes by which accurate and false memories are produced. The same inferential process is at work in *all* attempts at remembering; “veridical memories ... are no less constructed than false memories” (Sutton & Windhorst, 2009, p. 87).

Put another way, Constructivists collapse the processing distinction between memory errors and successful remembering. Memory errors are not malfunctions. Memory errors and successful remembering can, of course, be distinguished by their accuracy, at least in cases where external confirmation is available. But this does not reflect a substantive distinction between the processes by which such representations are generated. Rather, it is indicative of the fact that memory’s reliability depends on the similarity between what one wants to recall and the information available for constructing a memory. As De Brigard explains, “most of the time what you recall accurately depicts the witnessed event. Sometimes it does not. In both cases, however, the system is doing what it is supposed to do” (2014, p. 172).

Constructivists argue that the processing shared by successes and errors explains the finding that confidence and accuracy in remembering are orthogonal, as discussed in section 2.2. Given that accurate and false memories are produced by the same process, the Constructivist argues, it makes sense that participants are often unable to detect a difference between them. Moreover, Constructivists insist that this memory process, while sometimes leading to errors, is adaptive overall (Schacter & Addis, 2007). As evidence for this interpretation, they cite the fact that persons with severe memory disorders, such as amnesia, do not fall prey to memory manipulations like the DRM effect (Schacter, Verfaellie, & Koutstaal, 2002) and that susceptibility to DRM-style memory errors is, in fact, well associated with positive traits, like measures of creativity (Dewhurst, Thorley, Hammond, & Ornerod, 2011).

6. The Constructive View and the DRM

At a first pass, the Constructive View appears well suited for explaining the memory errors found in the DRM task. The “recognition” of *doctor* is a construction, built from—and distorted by—the similarities between the words on the list and the cue presented at recall. There are, however, important differences among versions of Constructivism in terms of how they characterize the constructive process by which representations are built. Below I consider each of the three to determine which, if any, best explains the DRM results.

6.1. Connectionism and the DRM Effect

For Connectionists, recall that the constructive process of remembering is guided by patterns of connectivity in the network of information that has been gathered from multiple past events. Answering recognition questions in the DRM task, is a matter of determining how easily a particular idea comes to mind (i.e., how quickly or fluently the node corresponding to the recognition cue or prompt is activated). Both *hospital* and *doctor* will be called to mind easily—*hospital* has been activated recently, and *doctor* has many nearby and closely associated nodes that have been activated recently. The similarity in the underlying activation of these two words explains the similarity of participant responses for these items. Words like *apple* that have not been activated recently will be less easily accessible and thus lack the feeling of familiarity, leading participants to deny recognizing them.

The Connectionist explanation of the DRM effect works well in cases where the recognition task follows soon after the presentation of the list, when the overall system is set in ways that, temporarily, reflect its influence. It is their temporal proximity that explains the influence of the list items on the activation of the node corresponding to the related but not presented word. If the recognition event occurs much after the initial event, then changes to the network during the interim should cause the

effect of this particular list-learning event to wash out. And, given the principles of spreading activation, once the event has receded into the general patterns of the network, it should not be recoverable.

The DRM effect can occur even when the learning event and the subsequent construction are separated in time. While standard versions of the task only separate the learning and recognition events by a few minutes, the effect remains when the events are separated by hours, days, and—in some cases—months (Seamon et al., 2002). Connectionists are at a loss to explain this. In order to be influenced by the general association between list items and non-presented lures, the participant must have a way of calling to mind the previous list. Otherwise those portions of the network are not active enough to support either accurate recognition or misremembering. Accommodating these results would require a way of tracking the influence of a particular past experience, but according to the principles upon which these underlying networks are built, there is no way for this to be done. So while it is true that it is the general relation between items on the list and the word *doctor*—a relation which predates the learning event—that explains the error in the DRM task, the ability to make use of that long-standing semantic relationship relies upon the retained influence of the particular past event. The ability to reactivate event-specific information is beyond the explanatory grip of the Connectionist. Any element of success due to a particular event, rather than a general pattern, is not amenable to the view.

6.2. Gist-Based Accounts and the DRM Effect

Gist-based accounts of remembering characterize the process as one of determining the degree of fit, or coherence, between the cue or prompt and the retained gist. Since all of the words on the initial list were *doctor*-related, this concept is embedded in the gist. This makes the word familiar to the participant, leading him or her to “recognize” the word as an item from the previous list. It also explains why *hospital* is recognized, but *apple* is not, as the former is more closely related to the event’s gist than the latter.

While Gist-based accounts appear to do well at explaining the DRM effect, full consideration of the range of DRM-style errors puts strain on this explanatory approach. First, as discussed in section 3.1, participants in DRM experiments not only falsely recognize words like *doctor*, their reports indicate that they feel that they remember these words. They deny that their responses are based on a general feeling of familiarity. Moreover, participants are often willing to provide various details about the experience of hearing or seeing the non-presented item. It is unclear why the retention of a non-detailed gist would encourage the generation of a detailed account of an event that did not happen. Second, DRM effects can be found for items that are similar to the initial event along many dimensions including not only semantic resemblance, as is the case in the *doctor* example, but also category membership and phonological or visual features. A participant in the DRM could schematize the experimental event as one of list-learning, or of volunteering, or of visiting a college campus, or of rainy Monday mornings. While each of these dimensions of similarity marks one way that an event’s gist could be retained, it is a stretch to characterize them all as dimensions of the *same* gist. To accommodate every possible dimension of similarity, proponents of the Gist-based view must allow the rememberer to retain several Gist-based representations of each past event, or allow for a Gist-based representation with enough detail to account for these various dimensions. Either way, it will lose its claim to being a less-demanding view of information processing and storage. Whatever success Gist-based accounts enjoy comes from their resemblance to, rather than distinctness from, the Archival View and its commitment to detailed representations of particular past events.

6.3. Episodic Hypothetical Thinking and the DRM Effect

According to De Brigard’s Episodic Hypothetical Thinking account, DRM errors are to be understood as instances of remembering what could have been the case. Recognition responses to *hospital* and *doctor* are similar because they have similar likelihoods for being on a list of *doctor*-related items. When

doctor is recognized, there is no malfunction because the response was produced by an (all things considered) reliable cognitive process. Claiming to recognize *doctor* is a good bet, one that—more often than not—pays off. Further evidence of the memory system's reliability comes from noting that less likely items, like *apple*, are not recognized.

This explanation of the DRM effect involves a critical, but often underemphasized, condition: the rememberer must retain information about the particular past event. Without appeal to such a discrete representation, there is no way to constrain one's consideration of which details were most likely part of that event. *Doctor* seems likely to have been on the list only if one's consideration of the possibilities is framed by memory of what other words were on the list (and further memory that the previous event one is being asked about involved list-learning). DRM results cannot be explained by frequency alone. While participant responses are influenced by overall frequency, DRM results cannot be explained by frequency alone, as explained in section 4.1. The influence of frequency is mediated by information from the list-learning event (Gallo & Roediger, 2003). That is, the recognition question is not one of unconstrained familiarity. The question is not, "is *doctor* a plausible word on a list, all things considered?" but rather, "is *doctor* a plausible candidate from the list of words I heard during the initial phase of the experiment?" Without appealing to a discrete representation of the particular past event of list-learning, the Episodic Hypothetical Thinking account cannot explain the DRM effect.

Although the Constructive View was designed to explain errors like the DRM, each particular version strains to accommodate the effect in its entirety. The DRM effect occurs across a wide range of stimuli, encoding conditions, and testing methods. Each version of Constructivism captures at least some of what happens in at least some instances, but no one view accounts for them all. This is, in part, because there are so many possible influences on the remembering process. Privileging one makes it difficult to capture errors that appear to derive from the others. But there is a further, more serious limitation that all versions of Constructivism share. False memories in the DRM are influenced by information retained from the initial learning event. False recognition of *doctor*, but not *apple*, reflects an influence of the words that were featured on the previous list, which cannot be explained away (at least not in all cases) by appeal to a general constructive processes. As this close analysis of DRM results shows, appeal to discrete representations is more critical to explanations of false memory than proponents of Constructivism have realized or been willing to acknowledge.

7. Explaining Misremembering

The DRM effect is a pervasive and persistent memory error, one that an account of memory must be able to explain. It is therefore troubling that neither the Archival nor the Constructive View can do so. In this final section, I offer a diagnosis and suggest a way forward.

First, the diagnosis. The Archival View acknowledges that—at least some of the time—we form and retain discrete representations of our past experiences. Such retention explains how details from particular events are produced in cases of successful remembering and distorted in cases of misremembering. The appeal to memory traces explains remembering, but the account of misremembering it offers is incomplete. On its own, the Archival View cannot account for DRM errors. Archivalists are at a loss to explain how one could retain information from a past event and yet fail to produce an accurate representation of that event when remembering. The difficulty comes from their inability to decouple the retention of a discrete representation from what is produced in retrieval.¹⁵ Proponents of the Archival View tend to think of retrieval as nothing more than a neutral probe for measuring the contents of the memory store.¹⁶

The Constructive View corrects for this, capturing the vagaries of the retrieval process by identifying the various influences on the representations constructed during remembering. This enables the Constructivist to explain why memory errors so often involve information from other events, the retrieval context, the rememberer's aims and assumptions, and so on. The aim for descriptive adequacy is an inspiring and promising feature of the approach. But the influences to which Constructivists appeal can only explain so much. A key influence on misremembering errors is left out: information

retained from the particular past event. In the Constructivist's hands, all representations of past events—whether completely accurate, wholly contrived, or somewhere in between—are, in some sense, confabulations. The view offers no means for differentiation between constructions that are constrained by information retained from a particular past event and those that are not. Capturing this difference would require appeal to the retention of discrete representations of particular past events, whose existence Constructivists have been at pains to deny. The problem with the Archival View is the account of *retrieval* it offers, not its commitment to *discrete representations of particular past events*. In rejecting both, the Constructivist has swung too far in the other direction.

The Archivalist and the Constructivist may face distinct challenges in trying to explain the DRM effect, but they derive from a shared assumption—namely, that the representation produced in remembering will be a direct reflection of what is stored in memory. Archivalists believe that memory produces veridical representations of past events by storing discrete representations of those events. Constructivists recognize that memories of past events contain information from several sources and conclude that this reflects the generalized organization of the underlying memory store. Misremembering errors, like those found in the DRM, challenge this shared assumption. Such errors rely on successful retention of the targeted event, implying a disconnect between what is retained in memory and the representation produced in the act of remembering.

The way forward is to create a hybrid theory, combining the complementary strengths and weaknesses of the Archival and Constructive views. There is too little space remaining to provide such a theory here. Instead, I will conclude by sketching two routes via which a hybrid theory could be developed, along with my reasons for favoring the latter.

One way to build a hybrid theory would involve further development of the most modest of the three versions of Constructivism outlined above: the Gist-based account. Gist-based Constructivism treats the memory trace as the locus of compromise. It retains the Archivalist commitment to the existence of memory traces, but argues for a Constructivist understanding of these traces as fluid, flexible, and abstract. Aside from its difficulties in explaining the DRM effect (identified in section 6.2), there are two general reasons to be dissatisfied with the explanatory framework such a hybrid would provide.

First, the Gist-based approach precludes any opportunity for a taxonomy of memory errors. According to this view, all misremembering—as well as successful remembering—involves a representation constructed on the basis of a gist trace. Sometimes the representations generated will include details that were part of the past event. Sometimes the details will be fabrications, albeit plausible ones. These outputs can be distinguished, provided a source of external verification is available for the past events in question. This allows the proponent of such a view to sort the various attempts at remembering on a continuum from most to least accurate. But the availability of such a continuum should not be confused for an explanation of differences between them. The Gist-based account cannot explain why any particular construction lands where it does. Each is built from the same Gist-based trace. The explanatory strategy inclines us away from any further investigation into the details of the remembering process, including the possibility of discovering differences in how various representations are produced and distorted.

Second, the gist traces that are central to this hybrid theory are vaguely defined, and it is unclear how an account of them can be made more precise. What is contained in the gist, exactly? The gist is intended to reflect a distillation of the past event, keeping only its essential features. But events have no such essence to distill—they can be conceptualized in any number of ways, and indeed, the various kinds of misremembering that are possible suggest that the abstraction occurs along many dimensions. In section 6.2, I raised this point as a challenge to the claim that Gist-based accounts offer more efficient information storage. Here I use it to highlight the explanatory limitations of this route to producing a hybrid theory. In order to accommodate the myriad summaries available for any past event within the same gist, it would have to be understood as whatever retained information from the past event would make possible all of the various ways an event can be reconstructed. But if this is so, then the account looks more like a redescription of the phenomena than an explanation.

For these reasons, I favor an alternative approach to creating a hybrid theory. Instead of proposing traces that blend Archival and Constructive features, as the Gist-based view does, I suggest an approach that incorporates aspects of the Archival and Constructive views as distinct components. Specifically, this approach would combine the Archival commitment to discrete retention with the Constructive approach to retrieval. Such a hybrid theory would include two separate conditions on successful remembering, one for each component: retention of information from a particular past event and construction of an accurate representation of that event at the time of retrieval. Each condition reflects a distinct influence on the remembering process that corresponds to the influences emphasized by each of the above perspectives. Disruptions of either will produce memory errors, but the errors will be different. The result is an explanation of misremembering, and its distinction from both successful remembering and other errors, such as confabulation. Misremembering occurs when the first retention condition is met, but the second accuracy condition is not. When neither condition is met, the result is confabulation.¹⁷ The approach has the ability to provide a taxonomy of memory errors, which I have argued is difficult for other hybrid approaches to supply. Moreover, attention to the various constructive processes and their respective influences on the accuracy condition may lead to an enhanced understanding of the mechanism by which rememberings and misrememberings are produced.

This is only a sketch. Determining which, if either, of these two approaches to a hybrid theory is preferable will require considerably more detail and development. Much of the requisite work will need to focus on the nature of memory traces. There is thus reason to expect that attention to work in neuroscience—on reconsolidation, binding, and the possibility of multiple trace formation—will provide an important guide.¹⁸ The need to account for the DRM effect, one of the most well-established effects in the psychological study of memory, should offer reason enough to undertake the effort.

Notes

1. For more details about the DRM paradigm, see Robins (2014).
2. When stimuli are lists of related words, backward associative strength is a strong predictor of the magnitude of false recognition (Cahn, McRae, & Katz, 2011).
3. Recall: Brainerd, Payne, Wright, and Reyna (2003). Pictures: Koustaal (2006). Faces: Homa et al., (2001). Dot arrays: Nosofsky (1991). The magnitude of the misremembering effect differs as a function of stimuli, which may warrant positing distinct mechanisms (e.g., Koutstaal & Schacter, 1997).
4. Children below the age of nine or ten are rarely susceptible to the DRM effect (e.g., Sugrue, Strange, & Hayne, 2009). I do not have the space to discuss the developmental trajectory of this error here, but doing so is an important further test of any theory's descriptive adequacy.
5. I do not intend to endorse any particular theory of perception, only the widely accepted characterization of perceptual errors (Fish, 2009).
6. I am grateful to an anonymous reviewer for this suggestion. In the future, it will be of interest to explore this issue and the errors near that fall near the alleged boundary. For now it is enough that there are clear cases on either side.
7. In referring to this view as archival, I am following Campbell (2006).
8. This view of retrieval is illustrated clearly in both Atkinson and Shiffrin's (1968) functional depiction of short and long-term memory and Anderson's ACT-R model of cognitive processing (Anderson, Bothell, Lebiere, & Matessa, 1998).
9. These four forms of error are intended to reflect the types of malfunction identified in section 3.3. This list may not be exhaustive.
10. The effect is well replicated (Gallo, 2006).
11. Support for Constructivism can also be found amongst philosophers who support theories of the mind that push back against traditional, faculty-based views of cognition and computational models of cognitive architecture (e.g., Bechtel & Abrahamsen, 2002; Dreyfus, 2002).
12. Gist-based views are also popular amongst memory scientists—see, for example, Brainerd and Reyna (2005) and Klein (2013).
13. Some call this construction, others *reconstruction*, depending on whether they care to emphasize the point that the initial encoding during perception is also constructive.

14. According to Fuzzy Trace Theory, there are two traces that are formed as the result of any given experience: a *verbatim* trace and a *gist* trace (Brainerd & Reyna, 2005). The verbatim trace—the detailed record of a past experience—decays rapidly and thus fails to play a role in most instances of remembering.
15. The Archivalist is not committed to the claim that everything is retained in memory, only that the information that is retained can be retrieved.
16. This view of retrieval is reflected in the methodological design of many memory experiments. Consider, for example, Ebbinghaus' (1913) diligent records of how many learning trials were required to master a list of syllables, or Craik and Tulving's (1975) exploration of which learning strategies lead to the best recall.
17. Another error—relearning—occurs when the second condition is met, but not the first. I have not discussed relearning in this paper, but it plays an important role in the Causal Theory of Memory (Bernecker, 2010; Martin & Deutscher, 1966). It is a virtue of the hybrid approach I propose that it can accommodate this error as well.
18. I am grateful to two anonymous reviewers for suggesting this opportunity for interface between these research programs.

Acknowledgements

Many thanks to Lauren Ashwell, Carl F. Craver, Corey J. Maley, Karen Neander, Katherine Rickus, Henry L. Roediger, and Shannon Spaulding for helpful comments on drafts of this article.

References

- Anastasi, J. S., Rhodes, M. G., & Burns, M. C. (2000). Distinguishing between memory illusions and actual memories using phenomenological measurements and explicit warnings. *American Journal of Psychology*, *113*, 1–26.
- Anderson, J. R., Bothell, D., LeBièvre, C., & Matessa, M. (1998). An integrated theory of list memory. *Journal of Memory and Language*, *38*, 341–380.
- Aristotle. (1955). *Parva naturalia* (W. D. Ross, Trans.). Oxford: Clarendon.
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The Psychology of learning and motivation* (pp. 89–195). New York: Academic Press.
- Bartlett, F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press.
- Bechtel, W., & Abrahamsen, A. (2002). *Connectionism and the mind: An introduction to parallel processing in networks* (2nd ed.). Cambridge, MA: Blackwell.
- Bernecker, S. (2010). *Memory: A philosophical study*. Oxford: Oxford University Press.
- Brainerd, C. J., & Reyna, V. F. (2002). Fuzzy-trace theory and false memory. *Current Directions in Psychological Science*, *11*, 164–169.
- Brainerd, C. J., Payne, D. G., Wright, R., & Reyna, V. F. (2003). Phantom recall. *Journal of Memory and Language*, *48*, 445–467.
- Brainerd, C. J., & Reyna, V. F. (2005). *The science of false memory*. Oxford: Oxford University Press.
- Cahn, D. R., McRae, K., & Katz, A. N. (2011). False recall in the Deese-Roediger-McDermott paradigm: the role of gist and associative strength. *Quarterly Journal of Experimental Psychology*, *64*, 1515–1542.
- Campbell, S. (2003). *Relational remembering: Rethinking the memory wars*. New York: Rowman & Littlefield.
- Craik, F. I. M., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*, 268–294.
- Debus, D. (2010). Accounting for epistemic relevance: A new problem for the causal theory of memory. *American Philosophical Quarterly*, *47*, 17–29.
- De Brigard, F. (2014). Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, *191*, 1–31.
- De Brigard, F., Addis, D., Ford, J. H., Schacter, D. L., & Giovanello, K. S. (2013). Remembering what could have happened: Neural correlates of episodic counterfactual thinking. *Neuropsychologia*, *51*(12), 2401–2414.
- Deese, J. (1959). Influence of inter-item associative strength upon immediate free recall. *Journal of Experimental Psychology*, *58*, 17–22.
- Dewhurst, S., & Farrand, P. (2004). Investigating the phenomenological characteristics of false recognition for categorized words. *European Journal of Cognitive Psychology*, *16*, 403–416.
- Dewhurst, S. A., Thorley, C., Hammond, E. R., & Ornerod, T. C. (2011). Convergent, but not divergent, thinking predicts susceptibility to associative memory illusions. *Personality and Individual Differences*, *51*, 73–76.
- Dreyfus, H. L. (2002). Intelligence without representation: Merleau-Ponty's critique of mental representation and the relevance of phenomenology to scientific explanation. *Phenomenology and the Cognitive Sciences*, *1*, 367–383.
- Ebbinghaus, H. (1913). *Memory: A contribution to experimental psychology* (H.A. Ruger & C.E. Bussenius, Trans.). New York: Teachers College. (Original work published 1885).
- Fish, W. (2009). *Perception, hallucination, and illusion*. Oxford: Oxford University Press.

- Gallo, D. A. (2006). *Associative illusions of memory: False memory research in DRM and related tasks*. New York: Taylor & Francis.
- Gallo, D. A., & Roediger, H. L. (2003). The effects of associations and aging on illusory recollection. *Memory & Cognition*, 31, 1036–1044.
- Homa, D., Smith, C., Macak, C., Johovich, J., & Osorio, D. (2001). Recognition of facial prototypes: The importance of categorical structure and degree of learning. *Journal of Memory & Language*, 44, 443–474.
- James, W. (1890). *The principles of psychology*. London: Macmillan.
- Klein, S. (2013). The temporal orientation of memory: It's time for a change of direction. *Journal of Research in Applied Memory and Cognition*, 2, 222–234.
- Koutstaal, W. (2006). Flexible remembering. *Psychonomic Bulletin & Review*, 13, 84–91.
- Koutstaal, W., & Schacter, D. (1997). Gist-based false recognition of pictures in older and younger adults. *Memory and Language*, 37, 555–583.
- Loftus, E. F. (2003). Our changeable memories: Legal and practical implications. *Nature Reviews: Neuroscience*, 4, 231–234.
- Loftus, E. F., & Pickrell, J. E. (1995). The formation of false memories. *Psychiatric Annals*, 25, 720–725.
- Martin, C. B., & Deutscher, M. (1966). Remembering. *Philosophical Review*, 75, 161–196.
- McDermott, K. B., & Roediger, H. L. (1998). Attempting to avoid illusory memories: Robust false recognition of associates persists under conditions of explicit warnings and immediate testing. *Journal of Memory & Language*, 39, 508–520.
- Michaelian, K. (2012). Generative memory. *Philosophical Psychology*, 24, 323–342.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton.
- Neisser, U. (1981). John Dean's memory: A case study. *Cognition*, 9, 102–115.
- Neisser, U., & Harsch, N. (1992). Phantom flashbulbs: False recollections of hearing the news about Challenger. In E. Winograd & U. Neisser (Eds.), *Affect and accuracy in recall: Studies of flashbulb memories* (pp. 9–31). Cambridge: Cambridge University Press.
- Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception & Performance*, 17, 3–27.
- Plaut, D. C. (1995). Semantic and associative priming in a distributed attractor network. Proceedings of the 17th Annual Conference of the Cognitive Science Society (pp. 37–42). Mahwah, NJ: Erlbaum.
- Robins, S. K. (2014). Memory traces, memory errors, and the possibility of neural lie detection. In C. T. Wolfe (Ed.), *Brain theory: Essays in critical NeuroPhilosophy* (pp. 171–191). New York: Palgrave.
- Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words that were not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 803–814.
- Ross, M. (1989). Relation of implicit theories to the construction of personal histories. *Psychological Review*, 96, 341–357.
- Schacter, D. L., & Addis, D. R. (2007). On the constructive episodic simulation of past and future events. *Behavioral & Brain Sciences*, 30, 299–351.
- Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology*, 49, 289–318.
- Schacter, D. L., Verfaellie, M., & Koutstaal, W. (2002). Memory illusions in amnesic patients: Findings and implications. In L. R. Squire & D. L. Schacter (Eds.), *Neuropsychology of memory* (3rd ed.) (pp. 114–129). New York: Guilford.
- Seamon, J. G., Luo, C. R., Kopecky, J. J., Price, C. A., Rothschild, L., Fung, N. S., & Schwartz, M. A. (2002). Are false memories more difficult to forget than accurate memories? The effect of retention interval on recall and recognition. *Memory & Cognition*, 30, 1054–1064.
- Sugrue, K., Strange, D., & Hayne, H. (2009). False memories in the DRM paradigm: Age-related differences in lure activation and source monitoring. *Experimental Psychology*, 56, 354–360.
- Sutton, J. (2007). *Philosophy and memory traces: Descartes to connectionism*. Cambridge: Cambridge University Press.
- Sutton, J., & Windhorst, C. (2009). Extended and constructive remembering: Two notes on Martin and Deutscher. *Crossroads: An Interdisciplinary Journal for the Study of History. Philosophy, Religion, and Classics*, 4, 79–91.