

MIT Press • Open Encyclopedia of Cognitive Science

Computational Models of Memory

Gregory E. Cox¹ Richard M. Shiffrin²

¹University at Albany, State University of New York, ²Indiana University

MIT Press

Published on: Jun 17, 2026

URL: <https://oecs.mit.edu/pub/8c02n2f1>

License: [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License \(CC-BY-NC-ND 4.0\)](#)

Our memories make us human, governing everything we are, do, act, decide, plan, and remember. Given the complexity of the mind and brain, simplified models are an essential tool for understanding how memory works. These models help researchers describe how experiences are represented in memory and the processes by which they are stored and retrieved. Although some models can be described using words, models expressed in formal terms, via mathematics or computer simulation, make explicit the basis for their explanations and allow precise predictions to be derived and tested. Computational models of memory are as diverse as the phenomena they address but tend to share a number of essential elements. They distinguish between memory for events with specific personal context (*episodic memories*) and memory for broadly relevant knowledge and skills (*semantic memories*); they distinguish between memory representations held in a temporarily active state (*short-term memories*) and those that are not currently active but could be (*long-term memories*); and they describe retrieval as a process that involves using a *probe* of memory, formed using available cues, to activate memory representations based on their degree of similarity to the probe.

History

The conceptual roots of this area stretch back to Aristotle, who proposed the idea that memories are evoked when someone experiences something similar to what they experienced at the time the memory was formed [see [Memory](#)]. It was not until the cognitive revolution beginning in the 1950s that these ideas were developed into the first formal (mathematical) models of memory and learning by researchers like William K. Estes, Richard C. Atkinson, Gordon Bower, and Robert R. Bush. Although memories cannot be observed directly, models made it possible to derive and test predictions about how different kinds of learning would result in differences in behavior, helping researchers to reverse engineer how memory works. These early models were comparatively simple, aiming to describe a small set of phenomena in terms of a few processes. As subsequent models aimed to capture a wider range of phenomena and to describe more processes in greater detail, purely mathematical derivations became increasingly difficult. It was not long before memory models were expressed primarily in the form of computer code with predictions obtained by simulation.

Core concepts

There is tremendous variety among computational models of memory, but all models are motivated by the desire to understand how particular phenomena arise as a function of how memories are stored and retrieved. Some models aim only to describe one or more phenomena; for example, a model could describe forgetting over time in terms of a mathematical function ([Averell & Heathcote, 2011](#)). By contrast, *causal* models aim to describe the processes responsible for producing an observed memory phenomenon. Such a model may explain how forgetting is caused by gradual degradation of memory representations, by interference from intervening events, by changes over time in the retrieval cues used to access memories, or some combination of these or other mechanisms.

Measuring memory

Although many memory phenomena are observed in daily life, models are particularly well suited to address quantitative data measured in controlled experiments. Some of the prominent experimental paradigms to which models are applied are listed in [Table 1](#), along with the kinds of data that are often obtained from those paradigms. As [Table 1](#) illustrates, knowledge and event memory, also referred to as semantic and episodic memory respectively ([Tulving, 1972](#)), tend to be investigated using different paradigms, and as such, these two types of memory are not often addressed by the same model. Even so, because these paradigms collect similar kinds of data (e.g., accuracy, confidence, and response time), models of event memory and models of knowledge often share many characteristics.

Table 1. *Prominent Experimental Paradigms for Studying Event Memory and Knowledge and Examples of the Kinds of Data Obtained*

Table 1			
	Paradigm	Description	Examples of types of data obtained

	designated item).	Degree of precision (size of error) in reproduction	
Knowledge or semantic memory	Lexical decision	Participants are presented with strings of letters and asked whether they form a known word or not.	Proportion of words correctly recognized (hit rate), proportion of nonwords falsely recognized (false alarm rate)
	Free association	Participants are provided with a cue word and asked to report any other words that this word brings to mind.	Total number of words reported and in what order
	Semantic fluency	Participants are asked to name as many members as they can from a known category (e.g., farm animals or words beginning with “s”).	

Representing information in memory

A representation of information in memory is termed a memory *trace* (also called an *engram*). A memory trace formed from an event contains information about the *content* of the event (what happened) as well as information about the *context* in which the event took place (e.g., where, when, with whom, and why it happened). Event traces encode information that was personally experienced. Knowledge traces differ from event traces in that they contain nonspecific context information because they represent information that has been encountered across a variety of contexts. For example, a knowledge trace may contain the content information that Laos is a landlocked country, whereas an event trace might represent the specific geography class period (the context) in which a person first encountered that information.

Activating memory traces

Memory traces can be in different states of *activation*, a term which refers to the ease with which information in the trace can be extracted and used. In most models, traces that are sufficiently active are termed *short-term* traces, in contrast to *long-term* traces that may or may not be active at any given time. The heightened activation of a short-term trace may last anywhere from fractions of a second to minutes, but it is temporary (and hence *short term*) because activation must be maintained either through deliberate effort, like rehearsal, or based on cues in the environment (e.g., if something is written on a post-it note nearby). The need for maintenance means that short-term memory has *limited capacity*; there is a limit to how much information can be maintained in short-term traces at any given moment without loss ([Oberauer et al., 2016](#)). Models distinguish many short-term memories, varying in the modality of the sensory input (e.g., auditory, visual, or haptic), capacity, and duration. For example, sensory short-term

traces have the capacity to represent many visual details, but only for a few hundred milliseconds, although it is possible to maintain a smaller number of less-detailed traces for a longer duration ([Sperling, 1960](#)) [see [Visual Memory](#)].

Short-term traces are particularly important because they are the basis for many of the processes that control storage and retrieval; these *control processes* are said to reside in a type of short-term memory termed *working memory* ([Atkinson & Shiffrin, 1968](#); [Baddeley & Hitch, 1974](#); [Wixted, 2024](#)) [see [Working Memory](#)]. Long-term traces act as a relatively permanent storehouse of information about both events and knowledge. New long-term traces are formed when information is successfully transferred from a short-term trace. As described below, traces in long-term memory may be activated during the process of retrieval, potentially reaching high enough activation to (re)enter short-term memory. However, successful retrieval is not guaranteed; long-term memory can be thought of as a gigantic library that does not have a precise index that would direct someone to a desired trace. Moreover, the information contained in a long-term trace is not necessarily immutable; long-term traces can be altered, augmented, and changed once retrieved. A negative consequence of this is that long-term traces may represent events in a distorted manner (e.g., [Loftus, 2005](#)). A positive consequence is that traces can accumulate information across events, enabling event memory and knowledge traces to coevolve with one another ([Nelson & Shiffrin, 2013](#)).

Storing information in memory traces

When a new event occurs, it is initially represented as a short-term memory trace, containing information retrieved from one or more long-term knowledge traces. For example, a short-term trace representing an encounter with the letter “Q” will contain more than just low-level sensory information such as edges, lines, and curves; it will also contain information from knowledge about the letter’s phonology and manner of usage, which is stored in knowledge traces that have accumulated information over a long history of language experience. As another example, watching a soccer game may activate knowledge about the players and teams as well as the rules of the game. As a final example, the experience of engaging in a conversation activates knowledge about the semantics and syntax of the language being used. The resulting long-term event traces formed from the kinds of experiences described in the preceding examples would thus include information derived from other knowledge traces.

Knowledge traces are augmented with new information encountered in the event, thus explaining how knowledge forms. For example, seeing a soccer player pull off a particular shot may result in the knowledge trace for that player being updated to reflect this ability, or hearing a word used in a new way may result in its knowledge trace being updated to reflect this additional shade of meaning. The relationships between the different types of information posited by memory models are depicted schematically in [Figure 1](#).




Figure 1

Schematic of the relationships between different representations posited by computational models of memory. Current experience is represented as a short-term trace containing both content and context information, with some content information derived from knowledge. The short-term trace is later transferred to a long-term trace in event (or episodic) memory and may also augment one or more knowledge traces (sometimes said to be in semantic memory).

Probing memory

The processes by which information is extracted from memory traces and used to decide how to act in a particular context can be broadly divided into *control processes*, which select an overall strategy for retrieval and decision-making, and *automatic processes*, which implement the strategies defined by the control processes ([Atkinson & Shiffrin, 1968](#); [Shiffrin & Schneider, 1977](#)). In any given scenario, retrieval depends on the *cues* selected by the control processes, which are used to form a *probe* of memory. The cues that are selected could be external (e.g., the visual features of a person's face, which may cue retrieval of that person's name) or could be internal, including cues retrieved from memory (e.g., if you retrieve that you met a person at lunch last Tuesday, that may help to cue retrieval of their name). Among the most important control processes in retrieval is the ability to *attend selectively* to some kinds of information over others when selecting the cues used to construct a probe. By doing so, people can rely on other (potentially automatic) retrieval processes to activate memory traces that may be more relevant to their current goals ([Nosofsky, 1986](#)).

Retrieval cues are used to construct a short-term trace that acts as a probe to activate additional memory representations. Different models represent the activation process in different ways, although all are based on a similarity of retrieval cues to traces ([Clark & Gronlund, 1996](#)). Some models represent memory traces as associative links between nodes in a network (e.g., [Gillund & Shiffrin, 1984](#); [Howard & Kahana, 2002](#); [Reder et al., 2000](#)). In these network models, activation *spreads* from the nodes corresponding to the

short-term trace of the probe to other nodes in proportion to the strength of the links connecting them. Other models represent memory traces as patterns of *feature values*; some models assume that traces of different events or knowledge are stored separately (e.g., [Cox, 2024](#); [Hintzman, 1988](#); [Shiffrin & Steyvers, 1997](#)), whereas others assume these traces are collapsed into a single composite structure (e.g., [Anderson, 1973](#); [Eich, 1982](#); [Murdoch, 1982](#)). Among each of these featural models, probes also consist of patterns of features, and traces are activated in proportion to their *similarity* to the probe (i.e., the degree to which the features in a trace match the corresponding features in the probe).

To illustrate how activation works in these kinds of models, imagine that you met someone at breakfast earlier in the day. The long-term memory trace for that event would contain information about its content, particularly the person's name and appearance, as well as the breakfast context. In a network model, this trace would be represented as a set of links connecting nodes, in which some nodes correspond to features of the person's appearance, some to the sound and/or spelling of their name, and some to features of the surrounding breakfast context. In feature-based models, those feature values would be *concatenated* (joined together) into a *vector* instead (see the section “Making decisions based on retrieved information”). When you see that same person later at lunch, you may attempt to retrieve their name. In a network model, this would amount to activating the nodes corresponding to the person's appearance as well as the features of the lunch context (which may or may not be similar to those of the earlier breakfast context). Activation would then spread across the links representing your earlier memory trace to (hopefully) activate features of the person's name. In a featural model, the probe would consist of a vector containing features of the person's appearance and the lunch context and would activate long-term traces that contained similar features, including (hopefully) the trace corresponding to your initial breakfast meeting. Although network and feature-based models can often be used to understand similar phenomena, feature-based models are well suited to address situations in which events may vary in their degree of similarity (e.g., two people may have the same eye color but different hair color).

Making decisions based on retrieved information

Some decisions, like those in lexical decision or item recognition, can be made on the basis of how much a probe activates the traces in memory. Many models applied to these kinds of *recognition* tasks treat the sum or average activation value across traces as a quantity called *familiarity* or *memory strength*, which is compared against a criterion to decide whether or not an event had been experienced (item recognition) or a word is known (lexical decision). Trace activation is a dynamic property that evolves as the probe changes over time, helping to account for the speed of retrieval ([Cox & Shiffrin, 2017](#)). To continue the breakfast example (see the section “Probing memory”), familiarity may be sufficient to decide whether or not one has recognized having met a person at breakfast, and one may be able to do so quickly if the probe strongly activates the trace formed at breakfast.

Other decisions, like those in cued or free recall or in free association or semantic fluency, require not just activating memory traces but extracting specific information from them. This is usually modeled as a *search* process; a candidate memory trace is *sampled* in proportion to its activation, and then relevant information has a chance to be *recovered* from the sampled trace. If the recovered information is deemed sufficient, a decision is made; otherwise, another sample may be taken, and the search continues, or a decision may be made to stop searching ([Raaijmakers & Shiffrin, 1980](#)). In the breakfast example, if the features of the person's name can be recovered well enough from the trace of your prior meeting, you may be able to greet them at lunch. Search can also be used for recognition when *recollection* of information from a strongly activated trace provides information for making a recognition decision; for example, you may recognize having seen someone earlier by recovering a particularly distinctive feature of that person from a memory trace.

An example of mathematical and computational modeling of memory

A concrete example of a specific computational model, retrieving effectively from memory (REM; [Shiffrin & Steyvers, 1997](#)), is provided in the context of the event memory task of item recognition (see [Table 1](#)). REM was designed to help understand how a number of important phenomena in item recognition could be produced by a single set of memory representations and retrieval processes. These phenomena include the fact that recognition accuracy diminishes with the number of items to remember (i.e., list length), that recognition is better for low-frequency words than high-frequency words, and that improving memory for some items on a list (e.g., by repetition) does not impair memory for other items on the list.

In a typical version of an item recognition task, a participant is shown a list of randomly selected words, one at a time, during a study phase. The presentation of each word constitutes a distinct event. The semantic and perceptual features of the word constitute the content for each event, whereas the context of each event is defined by the time, location, task goals, and internal states (like mood or motivation) of the participant. A distractor task, like solving a math problem or doing a puzzle, follows the study phase. The effect of the distractor is to deactivate any short-term traces of the studied items so that recognition is based only on long-term traces. After the distractor task, a test phase presents the participant with another list of words, also presented one at a time. Half of the words on the test list previously appeared on the study list (termed *targets*), whereas the other half had not been seen during the study phase (termed *foils*).

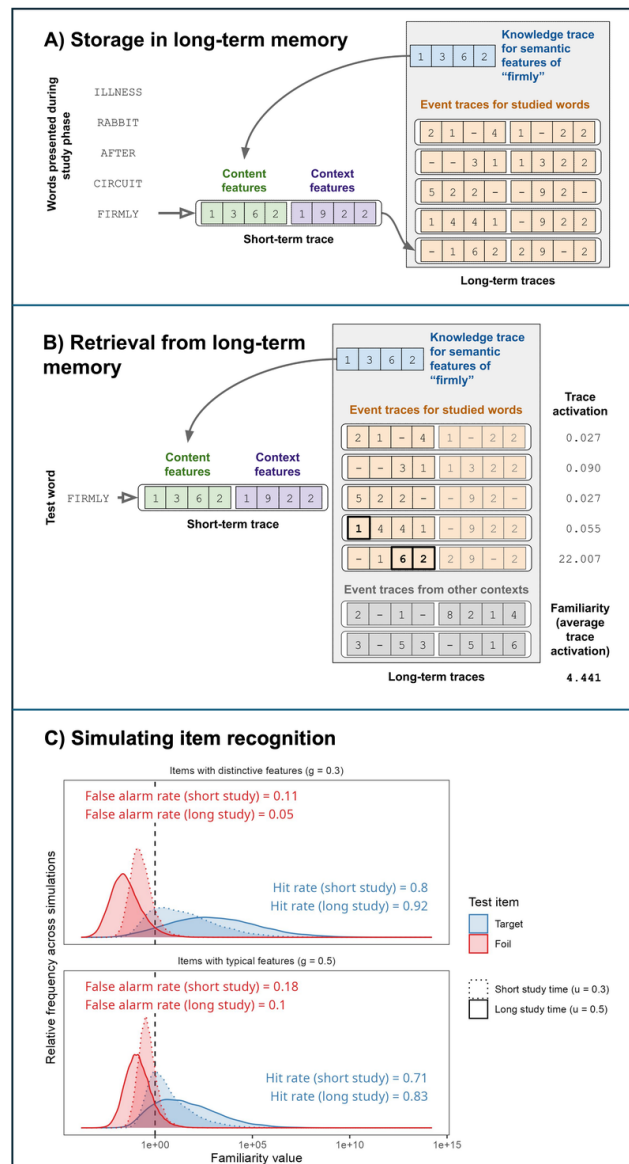


Figure 2

Example of the REM model applied to an item recognition task. Panel A illustrates how items from a study list end up represented as event traces in long-term memory. Panel B illustrates how a test item is used as a probe to activate event traces in long-term memory; context features deactivate any traces from contexts other than the list, whereas the remaining traces are activated in proportion to how many content features they share with those in the probe. Panel C shows the results of 20,000 simulations of REM performing item recognition using different parameters to model variations in study time (u parameter) and item distinctiveness (g parameter); in these simulations, the c parameter was fixed at 0.7, and it was assumed that 20 items had been studied (only 5 are shown in panels A and B for simplicity) and that each trace contained 20 content and 20 context features (only 4 of each are shown in panels A and B for simplicity).

In REM, each study event is represented as a separate long-term trace in event memory. Each of these traces is modeled as a vector of numbers, with the numbers representing the values of different features

of the content and context of the event. For example, one content feature may represent the first letter of the word, whereas a different content feature may represent some aspect of its meaning (the value of which is therefore retrieved from a knowledge trace). REM treats these feature values in a simplified abstract way, such that smaller numbers correspond to feature values that occur more frequently in the environment, whereas larger numbers correspond to rare feature values. For example, a feature representing the first letter of a word might have a low number for “shoe” (since “s” is a common letter) but a high number for “zebra” (since “z” is a rare letter). The relative frequency of common values over rare values is described by a parameter g . Because smaller values of g allow for rare features to occur more often, this parameter can be used to model situations in which some items are more distinctive than others (e.g., low frequency words also tend to have uncommon spellings and meaning).

Table 2

Parameters of the REM Model of Event Memory

Table 2	
Parameter	Interpretation
g	Degree to which items tend to either share many common features (values closer to one) or have distinctive features (values closer to zero)
u	Probability of transferring a feature of an event from a temporary short-term trace into a long-term trace
c	Probability that a feature transferred into a long-term trace is stored correctly

When a word is studied, only some of its features are correctly transferred from its initial short-term trace into its long-term trace. The probability of transferring a feature at all is u , and the probability of transferring it correctly is c . If a feature is transferred incorrectly, a random value is transferred instead, sampled in proportion to the value’s base rate frequency (described by parameter g). The result is an incomplete and error-prone memory trace consisting of values that describe certain features of the studied word. The u parameter can be used to model situations in which some items are studied more often or for a longer time, thereby increasing the amount of information about them that gets retained in long-term memory. The c parameter may be used to model other sources of variability, such as the ability to correctly perceive a feature (e.g., the c parameter may be lower if a word is printed in a hard-to-read font). These parameters and their interpretation are summarized in [Table 2](#).

To make a decision about whether a word shown during the test phase was or was not presented during the study phase, the short-term trace representing the test word is treated as a probe that is compared to each of the long-term traces in memory. First, the context features in the probe are used to filter out any

traces from contexts other than the study phase. For the remaining traces, each content feature that matches the probe increases the trace's activation, whereas each mismatch decreases the trace's activation. As a result, if the test word is a target, the long-term trace of when the target had been studied will tend to be strongly activated unless it happened to have been stored poorly. Conversely, if the test word is a foil, most event traces would have low activation unless they happened to share features with the test word by chance, or errors in storage resulted in a trace with similar features. The average of the trace activation values is termed *familiarity*. REM predicts that a participant recognizes an item when familiarity exceeds a criterion value of one; otherwise, they will decide the test word was not on the study list.

REM is an example of how a computational model explains overt actions (responses in an item recognition task) in terms of memory representations (traces that record the features of past experiences) and retrieval processes (activation of those traces by a probe). Because the features of the study and test words can vary and because there is noise in how well those features are transferred to long-term memory, REM predicts an entire distribution of possible familiarity values that a target or foil might produce. By simulating REM performing item recognition many times, it is possible to estimate these distributions and how they change as a function of REM's parameters, as illustrated in [Figure 2C](#). By computing the proportion of simulations in which familiarity exceeds REM's criterion, REM can be used to predict the probability that a participant will either correctly recognize a target (hit rate) or falsely recognize a foil (false alarm rate) [see [Signal Detection Theory](#)].

Model complexity

Although REM accounts for some important phenomena in item recognition, it is missing many other important processes and hardly accounts for all possible phenomena. For example, it does not model the alteration of existing traces or the formation of new memory traces during the test phase nor does it account for the speed or confidence of responses. No computational model can account for everything, nor should it. A complete model of how a specific person remembered a specific piece of information in a specific context would require a model of that person's entire life history, the social and natural environment in which they lived, and a complete account of their thinking processes and neurophysiology. Even if all of that were somehow able to be modeled, the result would be so complex that no individual person could understand it, and it would be so specialized that it would not be able to explain or predict anything beyond its narrow domain. The value of a model like REM is that it helps researchers understand how a selection of clearly defined phenomena arise from the operation of the processes described by the model.

Any model is, therefore, an intentional simplification of reality. Memory models—indeed, all models of anything—are “wrong” in that they fail to capture every aspect of their domain of application. Even so, simple models are useful when they approximate the most important processes responsible for producing replicable phenomena that generalize to a variety of settings. Moreover, simpler models are

easier to understand and describe, helping scientists understand how predictions follow from the assumptions of the model. With simpler models, it is easier to explore the full range of a model's predictions across its parameter values, thereby helping to distinguish between qualitative patterns that the model always predicts from those that may only be predicted with a particular choice of parameter values. Finally, starting with a simple model makes it easier to use the model as an “intuition pump” to explore different variations of the model that add, remove, or alter some of its components. This last point is critical because building a model is an iterative process in which the failures of earlier versions suggest ways the model may be improved (for further discussion of these and related points, see [Cox & Shiffrin, 2024](#)).

Questions, controversies, and new developments

Given the diversity among memory models and paradigms, there remain many open questions regarding how best to use computational models to improve our understanding of memory. For one thing, memory models are beginning to address the fact that different people will often form memory traces containing different kinds of information, even from the same event ([Cox, 2025](#); [Nosofsky et al., 2025](#)). Researchers are also working to understand how the processes and representations described by computational models of memory are implemented in neural systems, particularly the hippocampus ([Norman & O'Reilly, 2003](#)), which appears to be critical for event memory, whereas knowledge depends more on cortical structures ([McClelland et al., 1995](#)). Relatedly, it remains unclear how content repeated across event traces evolves into knowledge traces, although models have begun to address this issue from various perspectives ([Jamieson et al., 2018](#); [Kumaran & McClelland, 2012](#); [Nelson & Shiffrin, 2013](#)). Work is also ongoing to understand the control processes that guide encoding and retrieval from memory, ranging from understanding the interplay between selective attention and memory ([Weichart et al., 2024](#)) to understanding how people optimize their use of memory ([Zhang et al., 2023](#)). Finally, the nature of capacity limits in short-term memories remains an active domain of current research (e.g., [Oberauer et al., 2016](#)).

Broader connections

Computational models are well suited to help researchers understand differences in memory between individuals, including differences because of aging ([Benjamin, 2010](#); [Healey & Kahana, 2016](#)) or cognitive impairments ([Nosofsky & Zaki, 1998](#)). Computational models of memory also help researchers understand the differences and similarities between biological memory and memory in artificial intelligence models ([Lu et al., 2022](#)). Memory models also help illustrate how aspects of judgment and decision-making can be attributed to the information in memory and how it is retrieved ([Thomas et al., 2008](#); [Wang et al., 2025](#)). Finally, computational models of memory can be useful in educational settings for helping to design instruction that fosters not only long-term retention but also the ability to build on existing knowledge to generalize and adapt to new settings ([Tabibian et al., 2019](#)).

Further reading

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *Psychology of learning and motivation* (pp. 89–195). Elsevier.
- Cox, G. E. (2024). Dynamic retrieval of events and associations from memory: An integrated account of item and associative recognition. *Psychological Review*, 131(6), 1297–1336.
<https://doi.org/10.1037/rev0000486>
- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3), 269–299. <https://doi.org/10.1006/jmps.2001.1388>
- Malmberg, K. J. (2024). *Human memory: The general theory and its various models*. Cambridge University Press.

References

- Anderson, J. A. (1973). A theory for the recognition of items from short memorized lists. *Psychological Review*, 80(6), 417–438. <https://doi.org/10.1037/h0035486>
[↩](#)
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *Psychology of learning and motivation* (pp. 89–195). Elsevier.
[↩](#)
- Averell, L., & Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1), 25–35. <https://doi.org/10.1016/j.jmp.2010.08.009>
[↩](#)
- Baddeley, A. D., & Hitch, G. (1974). Working memory. In G. W. Bower (Ed.), *Psychology of learning and motivation* (pp. 47–89). Elsevier.
[↩](#)
- Benjamin, A. S. (2010). Representational explanations of “process” dissociations in recognition: The DRYAD theory of aging and memory judgments. *Psychological Review*, 117(4), 1055–1079.
<https://doi.org/10.1037/a0020810>
[↩](#)
- Clark, S. E., & Gronlund, S. D. (1996). Global matching models of recognition memory: How the models match the data. *Psychonomic Bulletin & Review*, 3(1), 37–60. <https://doi.org/10.3758/bf03210740>
[↩](#)

- Cox, G. E. (2024). Dynamic retrieval of events and associations from memory: An integrated account of item and associative recognition. *Psychological Review*, 131(6), 1297–1336.
<https://doi.org/10.1037/rev0000486>

↩

- Cox, G. E. (2025). Similarity as likelihood ratio: Coupling representations from machine learning (and other sources) with cognitive models. *Psychonomic Bulletin & Review*, 33(1), 6.
<https://doi.org/10.3758/s13423-025-02828-w>

↩

- Cox, G. E., & Shiffrin, R. M. (2017). A dynamic approach to recognition memory. *Psychological Review*, 124(6), 795–860. <https://doi.org/10.1037/rev0000076>

↩

- Cox, G. E., & Shiffrin, R. M. (2024). Computational models of event memory. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford handbook of human memory* (pp. 64–93). Oxford University Press.

↩

- Eich, J. M. (1982). A composite holographic associative recall model. *Psychological Review*, 89(6), 627–661.
<https://doi.org/10.1037/0033-295x.89.6.627>

↩

- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91(1), 1–67. <https://doi.org/10.1037/0033-295x.91.1.1>

↩

- Healey, M. K., & Kahana, M. J. (2016). A four-component model of age-related memory change. *Psychological Review*, 123(1), 23–69. <https://doi.org/10.1037/rev0000015>

↩

- Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review*, 95(4), 528–551. <https://doi.org/10.1037/0033-295x.95.4.528>

↩

- Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3), 269–299. <https://doi.org/10.1006/jmps.2001.1388>

↩

- Jamieson, R. K., Avery, J. E., Johns, B. T., & Jones, M. N. (2018). An instance theory of semantic memory. *Computational Brain & Behavior*, 1(2), 119–136. <https://doi.org/10.1007/s42113-018-0008-2>

↩

- Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, 119(3), 573–616. <https://doi.org/10.1037/a0028681>

↩

- Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory*, 12(4), 361–366. <https://doi.org/10.1101/lm.94705>

↩

- Lu, Q., Hasson, U., & Norman, K. A. (2022). A neural network model of when to retrieve and encode episodic memories. *eLife*, 11, e74445. <https://doi.org/10.7554/elife.74445>

↩

- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3), 419–457. <https://doi.org/10.1037/0033-295x.102.3.419>

↩

- Murdock, B. B. (1982). A theory for the storage and retrieval of item and associative information. *Psychological Review*, 89(6), 609–626. <https://doi.org/10.1037/0033-295x.89.6.609>

↩

- Nelson, A. B., & Shiffrin, R. M. (2013). The co-evolution of knowledge and event memory. *Psychological Review*, 120(2), 356–394. <https://doi.org/10.1037/a0032020>

↩

- Norman, K. A., & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review*, 110(4), 611–646. <https://doi.org/10.1037/0033-295x.110.4.611>

↩

- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39–57. <https://doi.org/10.1037/0096-3445.115.1.39>

↩

- Nosofsky, R. M., & Zaki, S. R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. *Psychological Science*, 9(4), 247–255. <https://doi.org/10.1111/1467-9280.00051>

↩

- Nosofsky, R. M., Ehinger, K. A., & Osth, A. F. (2025). Tests of a hybrid-similarity exemplar model of context-dependent memorability in a high-dimensional real-world category domain. *Journal of Experimental Psychology: General*, 154(12), 3450–3480. <https://doi.org/10.1037/xge0001846>

↩

- Oberauer, K., Farrell, S., Jarrold, C., & Lewandowsky, S. (2016). What limits working memory capacity? *Psychological Bulletin*, 142(7), 758–799. <https://doi.org/10.1037/bul0000046>

↩

- Raaijmakers, J. G. W., & Shiffrin, R. M. (1980). SAM: A theory of probabilistic search of associative memory. In G. H. Bower (Ed.), *Psychology of learning and motivation* (pp. 207–262). Elsevier.

↩

- Reder, L. M., Nhouyvanisvong, A., Schunn, C. D., Ayers, M. S., Angstadt, P., & Hiraki, K. (2000). A mechanistic account of the mirror effect for word frequency: A computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(2), 294–320. <https://doi.org/10.1037/0278-7393.26.2.294>

↩

- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84(2), 127–190. <https://doi.org/10.1037/0033-295X.84.2.127>

↩

- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM—retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145–166. <https://doi.org/10.3758/bf03209391>

↩

- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11), 1–29. <https://doi.org/10.1037/h0093759>

↩

- Tabibian, B., Upadhyay, U., De, A., Zarezade, A., Schölkopf, B., & Gomez-Rodriguez, M. (2019). Enhancing human learning via spaced repetition optimization. *Proceedings of the National Academy of Sciences*, 116(10), 3988–3993. <https://doi.org/10.1073/pnas.1815156116>

↩

- Thomas, R. P., Dougherty, M. R., Sprenger, A. M., & Harbison, J. I. (2008). Diagnostic hypothesis generation and human judgment. *Psychological Review*, 115(1), 155–185. <https://doi.org/10.1037/0033-295x.115.1.155>

↩

- Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 382-402). Academic Press.

[↩](#)

- Wang, F., Aka, A., He, L., & Bhatia, S. (2025). Memory modeling of counterfactual generation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 51(2), 255–284.
<https://doi.org/10.1037/xlm0001335>

[↩](#)

- Weichart, E. R., Unger, L., King, N., Sloutsky, V. M., & Turner, B. M. (2024). “The eyes are the window to the representation”: Linking gaze to memory precision and decision weights in object discrimination tasks. *Psychological Review*, 131(4), 1045–1067. <https://doi.org/10.1037/rev0000475>

[↩](#)

- Wixted, J. T. (2024). Atkinson and Shiffrin’s (1968) influential model overshadowed their contemporary theory of human memory. *Journal of Memory and Language*, 136, 104471.
<https://doi.org/10.1016/j.jml.2023.104471>

[↩](#)

- Zhang, Q., Griffiths, T. L., & Norman, K. A. (2023). Optimal policies for free recall. *Psychological Review*, 130(4), 1104–1124. <https://doi.org/10.1037/rev0000375>

[↩](#)