

17 Storage Side Effects: Studying Processing to Understand Learning

Lawrence W. Barsalou

What does it mean to say that learning is goal-driven? As Ram and Leake outline in Chapter 1, goals could direct learning in several manners: behaviorally, strategically, or during storage. Behaviorally, goals could determine the information seeking activities that people pursue, such as consulting experts, reading texts, observing the environment and so forth. When people have the goal of purchasing a stereo, for example, it could lead them to consult knowledgeable friends, to read stereo magazines, or to observe stereo equipment. Strategically, goals could determine the cognitive procedures that people use to acquire, organize, and use information. Numerous strategies, such as rehearsal, elaboration, search, analogy, pattern completion, failure-driven learning, and so forth, could be accessed and deployed during goal pursuit. While stereo shopping, people could rehearse, elaborate, and organize information about a particular stereo receiver to remember and compare it to other receivers later; or people could combine utilities about a receiver's features to establish their preference for it relative to other receivers; or, based on case-based reasoning, people could draw inferences about a receiver's performance from memories of previous ones. Finally, goals could affect the storage of information during learning, with only goal-relevant information being stored and goal-irrelevant information being expunged (i.e., information filtering). While stereo shopping, people might only store information about receivers that are potential candidates and not store information about non-candidates; or people might only store features of receivers that are relevant to their goals and not store irrelevant features.

The propositions that goals constrain information-seeking behavior and cognitive strategies are noncontroversial. Clearly, the particular goals that people pursue affect their behavior, both physically and cognitively. Less obvious is the relation between goals and storage: Do people have the ability to "turn storage on and off," depending on the relevance of information to their goals? Or does the process of storage lie beyond strategic control?

As Ram and Leake's overview in Chapter 1 suggests, researchers who study goal-driven learning often believe that people have control over storage. According to this view, people intentionally store information that is relevant to their goals and intentionally discard information that is not. For

example, "learning goals ... affect learner decisions of when and what to learn" (p. 1); "the effectiveness of goal-driven learning depends on being able to make good decisions about when and what to learn" (p. 1); "[Goals] give criteria for evaluating the results of learning and deciding what learned information to store" (p. 3); "Evaluating the results of learning allows a system to decide whether to store them, and, if it does, to decide how to store them" (p. 25); "Information filters can be used to decide which learned items to store, which to retain in memory over time" (p. 26).

Although I agree that goals almost certainly control people's actions and their cognitive strategies, I shall argue that goals do not affect storage directly: People do not have the ability to turn storage on and off at will as they construe passing information to be goal relevant or goal irrelevant. Instead, I shall argue that information storage is an automatic, uncontrollable side effect of processing, lying beyond people's strategic assessments of goal relevance and their attempts to control storage.

17.1 INTENTIONS AND INFORMATION STORAGE

Considerable evidence, and considerable agreement, exist in the human cognitive literature on the role of intentions in storage. First, it is well known that intentions are not necessary for storage, as the well-known phenomenon of incidental learning illustrates. In many memory experiments (e.g., Craik & Tulving, 1975; Hyde & Jenkins, 1969), and typically in real life, people are not asked to store their experiences, nor do they typically have the intention of storing them. Nevertheless, people almost always remember a considerable amount of information from these events, storing most of it automatically, without the intention of doing so. Incidental learning offers an existence proof that the intention to store information is not necessary for robust storage to occur.

What about having the intention to store information? Do such intentions change or improve information storage? Ultimately, the answer is no: Intentions have no direct effect on storage. However, intentions may affect storage indirectly, when they cause people to change information processing, which then affects storage directly.

Consider a concrete example. Eagle and Leiter (1964) gave subjects a list of words, asking an intentional group explicitly to remember them, and asking an incidental group only to classify each word as a part of speech. Because the incidental subjects were not warned about a later test, and because of the cover story about categorizing parts of speech, these subjects did not expect to be tested, and, thus, did not develop intentions to remember the words. Instead, these subjects simply believed that the experimenter was studying their ability to classify this material.

On a later recognition test, subjects had to discriminate presented words from nonpresented words. As indicated by higher discrimination accuracy, the incidental group actually remembered more than the intentional group. This

result demonstrates that the intention to store information is not necessary for significant information storage to occur. On a sensitive test of memory, the incidental subjects remembered a considerable number of the words that they had classified as parts of speech, even though they did not have the intention to store them.

In contrast, the intentional group performed better than the incidental group when recall memory was tested (i.e., having to reproduce the words seen earlier). The intention to store information could have been responsible for this effect. However, a much more compelling, and more widely accepted, account exists: The intention to store information caused subjects to deploy processes that they would not otherwise use, which had the side effect of increasing the storage of information useful for a later recall test.

Many subsequent experiments support the latter explanation. Consider Mandler (1967), in which one group of subjects sorted a list of words into categories and another group just read them. Within each of these two groups, half of the subjects were asked to remember the words (intentional learning), and half received no such instructions (incidental learning). Three results on a subsequent recall test are of primary interest here: First, intention had no effect on the recall of those subjects asked to categorize the list: Incidental subjects who categorized the list recalled as much as intentional subjects who categorized it. Second, categorization instructions had a major effect on incidental subjects: Incidental subjects who categorized the list recalled it much better than incidental subjects who just read it. Third, categorization instructions had no effect on intentional subjects: Intentional subjects who only read the list did just as well as intentional subjects who categorized it.

Together these findings suggest strongly that the key factor in storage was having organized the material through sorting, not having had the explicit intention to store it. Incidental subjects who organized the list remembered as much as both intentional groups, whereas incidental subjects who only read the list remembered substantially less. Interestingly, intentional subjects who only read the list performed as well as subjects who organized it. This particular finding suggests that intentional subjects who only read the list actually organized it. Much additional work in the memory literature supports this conclusion: When subjects have the intention to store material, their organization of it at encoding surfaces clearly in their serially organized clusters of related words at recall (e.g., Bousfield, 1953; Cofer, 1967; Puff, 1970; also see Tulving, 1962, 1964, 1966).

It is now widely accepted that the intention to learn only improves recall performance because of the organizational strategies that it engenders. As theories of recall typically assume, organizational strategies establish pathways in memory between presented words that later facilitate search and retrieval (e.g., Anderson & Bower, 1972; Gillund & Shiffrin, 1984). In the absence of having an intention to learn, performing the same organizational processing produces exactly the same pathways in memory and the same high level of performance.

Thus, goals do not constrain storage by determining whether people decide to store or not store information that is goal-relevant. Intentions do not turn storage on and off. Instead, the performance of certain critical processing—not the intention to store information—is the key factor in storage. Because people do not have the ability to turn storage on and off at will, goal-driven learning simply does not work this way. Instead, it works analogously to recall: Having a goal changes the nature of processing, thereby producing storage side effects that may be beneficial.

17.2 THE EXTENT OF STORAGE SIDE EFFECTS

Many cognitive psychologists believe that storage is an obligatory side effect of processing. As cognitive processing proceeds, a record of processing is established automatically in memory, with the intention to turn storage on and off playing no role whatsoever. A key question arises from this view: What cognitive processes store information as a side effect?

One possibility is that only strategic, goal-driven processes produce storage side effects, such as the “classic problem solving” considered in chapter 1. More automatic processes, not related to strategic goals, do not. If so, then strategic, goal-driven behavior would constrain the storage of information in memory: Only information related to strategic goals is stored, such that goals constrain storage. Information related to more automatic processes is not stored.

Indeed, it is almost axiomatic in the memory literature that strategic processing produces records in memory, where *strategic processing* refers to processing that is both deliberate and conscious, and requires extensive use of limited attentional resources (see Barsalou, 1992, chap. 3, for a review). As the amount of strategic processing allocated to information increases, so does memory for it. This basic cognitive relationship manifests itself in many ways: Memory for information improves the longer it is processed, the more it is rehearsed, the better it is elaborated, the more it is organized with other information, and so forth (see Barsalou, 1992, chap. 6, for a review). Again, however, these benefits result, not from the intention to store information, but from increased amounts of processing. The more strategic processing performed, the more storage side effects accrue in memory. Because the pursuit of goals in classic problem solving requires strategic processing, we should certainly expect that traces of such processing are stored as side effects. Over the course of problem solving, information about the problem episode is stored in memory, with the amount of storage reflecting the amount of processing.

Do storage side effects only accrue for strategic activities that are deliberate, conscious, and effortful, such as the pursuit of goals in classic problem solving? As we shall see, the answer is no: Storage side effects occur for a wide variety of other processes that occur more automatically, unconsciously, and

effortlessly. The storage of information goes considerably beyond strategic, goal-driven processing.

Consider the simple reading of a word. Much evidence indicates that reading a word is a relatively automatic process. In the Stroop paradigm, for example, subjects are presented with a word like "red" printed in blue ink and simply have to state that the ink color is blue—they do not have to read and identify the word as "red" (Stroop, 1935; also see Logan, 1980; MacLeod & Dunbar, 1988). Interestingly, time to identify the ink color of "red" in blue ink is slower than for a comparable noncolor word in blue ink (e.g., "led"). Even though reading the word interferes with the task, subjects cannot prevent themselves from reading it, because reading occurs obligatorily on presentation of the word.

Do such relatively automatic processes produce storage side effects? Much work in the lexical priming and implicit memory literatures indicates clearly that they do. For example, the lexical priming literature has demonstrated that reading a word on one occasion decreases the time to read it later (e.g., Scarborough, Cortese, & Scarborough, 1977). Even though reading the word is a relatively automatic process, information storage occurs that speeds subsequent reading.

The implicit memory literature contains many other examples of such effects. For example, if subjects read a word in a particular font, they may not remember ever having seen the word, much less having seen the word in that font. Nevertheless, reading the same word later will be faster if it is presented in the same font instead of a new one (e.g., Jacoby & Hayman, 1987). Even though reading is a relatively automatic process, it produces storage side effects that contain the detailed and idiosyncratic features of words. Much other research in the implicit memory literature similarly indicates that a variety of simple, nonstrategic processes, such as reading, categorization, and retrieval, all produce storage side effects (e.g., Cohen & Squire, 1980; Jacoby & Brooks, 1984; Schacter, 1987; Squire, 1987). Storage side effects do not just occur for strategic, goal-driven processing, as in classic problem solving.

The power law of practice provides further evidence that storage side effects occur for relatively automatic processes (e.g., Anderson, 1983; Newell & Rosenbloom, 1981). Even after thousands and thousands of trials at the same task, performance continues to improve. Such improvement suggests that each additional performance of the task, even after it is overlearned and relatively automatic, produces a storage side effect. Although such side effects may simply involve the strengthening of procedures, they nevertheless occur automatically as side effects of processing.

As these examples demonstrate, a process does not have to be deliberate, conscious, and effortful to produce storage. Although strategically pursuing a goal is likely to produce storage, so is the performance of a more automatic and peripheral process. Learning is far from limited to the strategic, goal-driven processing of classic problem solving.

17.3 THE UBIQUITY OF GOAL-DRIVEN STORAGE

Thus far, we have established the following five points:

1. People do not have the ability to turn the storage of information on and off, depending on its relevance to their goals.
2. Storage results as an automatic, obligatory side effect of processing.
3. Storage side effects occur in the absence of an intention to learn or store information.
4. The amount of processing performed determines the amount of information stored.
5. Storage side effects occur, not only for strategic, goal-driven processing, but also for processing that is relatively automatic, unconscious, and effortless.

Based on point 5, one might conclude that all cognitive processes produce storage side effects. Indeed, it is difficult, if not impossible, to think of any process that does not produce a storage side effect of some sort.

If all cognitive processes produce storage side effects, then the relation between goals and storage could take either of the following two forms: First, the hypothesis that goals constrain storage is simply wrong, because many processes besides strategic goal-pursuit store information. Second, these more automatic, nonstrategic processes could be viewed as goal driven in some sense, such that all storage is goal driven. In this latter view, any operation of the cognitive system ultimately serves some goal. These goals may not be conscious or explicit, but may instead have been hardwired into the brain as "system goals" through evolution. For example, if processing a word's font serves the goal of recognizing the word, and if this processing has the side effect of storing the font, then storage of the font was goal driven. Similarly, if daydreaming serves system goals such as relaxation and stimulation, and if daydreaming produces memories, then, again, we have storage side effects that are ultimately goal driven. Because all cognitive processes produce storage side effects, and because all cognitive processes serve goals, all cognitive processes produce storage that is goal driven.

Thus, two additional points can be added to the five established earlier:

6. Any act of cognitive processing is goal-driven, not just the strategic performance of classic problem solving.
7. All acts of cognitive processing produce storage side effects that are goal driven.

Although these may be interesting conclusions, they leave the issue of goal-driven storage vacuous. If all storage is goal driven, then there is nothing special about goal-driven storage—it is ubiquitous. Moreover, it is incorrect to assume that some storage is goal driven and other storage is not. All processing is goal driven, all processing produces storage, and all storage is goal driven. Classic problem solving is not unique in exhibiting this property.

17.4 UNDERSTANDING GOAL-DRIVEN STORAGE THROUGH ANALYSIS OF GOAL-DRIVEN PROCESSING

There must be more to say about goal-driven storage than simply to note its ubiquity. In particular, this conclusion is unsatisfactory because it provides no detailed understanding of how goal pursuit establishes knowledge in memory through storage. However, the thesis that processing automatically produces storage side effects suggests a powerful strategy for exploring this issue: The best way to study goal-relevant knowledge in memory is to study the processing that produces it. Because storage is a side effect of processing, the structure of knowledge mirrors the structure of processing.

To illustrate this research strategy, the following two sections apply it to classic problem solving and implicit orientation. These two analyses are admittedly sketchy, but they serve to exemplify how analyses of processing can inform analyses of storage and the knowledge that it produces.

17.4.1 Processing and Storage in Classic Problem Solving

Assume, following Newell and Simon (1972), that problem solving includes four processes:

1. A person consciously selects a goal and intends to achieve it.
2. The environment is assessed to see what must be changed to achieve the goal.
3. A search through memory is performed to find knowledge that will produce a satisfactory change (e.g., rules, cases).
4. The knowledge is applied.

Subsequently, the cycle is repeated, beginning with another assessment of the environment, the execution of further knowledge, if necessary, and so on. This cycle continues until the problem is solved, or the problem solver gives up. It is probably fair to say that the assess-search-execute cycle in steps 2, 3, and 4 constitutes a reasonable characterization of processing in most theories of problem solving.

How might this processing cycle inform us about the storage of problem solving knowledge? Most theories assume that the series of assess-search-execute cycles that achieves a goal is stored as a side effect. In some theories, these side effects are stored as rules that each represent one assess-search-execute cycle (e.g., Anderson, 1983; Newell, 1990). In other theories, these side effects are stored as a case that contains a trace of the entire series of processing cycles (e.g., Hammond, 1990; Kolodner & Simpson, 1989). In either approach, the information stored in memory reflects processing directly. If the human cognitive system does in fact process problems in this manner, then it is probably no accident that the dominant approaches to learning in classic problem solving reflect this process as well.

The idea that processing determines storage further explains the development of the knowledge that underlies skill. In chunking theories of skill, such as ACT* and SOAR, the integration of rules in working memory produces more compact rules in long-term memory, which speed processing on later occasions to produce the power law of practice (Anderson, 1983; Newell, 1990). In exemplar theories of skill, increased processing of a task produces increased numbers of cases in memory, such that the speed of retrieving and applying any one case increases according to the power law of practice (Logan, 1988). In either theory, the stored knowledge that produces expert performance results directly as a side effect of processing in working memory.

Theories of problem solving typically do not assume that storage of problem-solving knowledge occurs intentionally. Rather than a problem solver believing consciously that a solution should be remembered, there typically is no such intention. Instead, the problem solver is simply trying to solve the problem, not thinking much about its future relevance, with storage resulting automatically. It is as if evolution decided it could not entrust storage to the intentions of problem solvers, opting instead to produce it automatically as a side effect.

Clearly, this type of storage benefits an intelligent system. Rather than having to solve every goal from scratch, the cognitive system stores information about previous problem-solving episodes to make future problem solving more efficient. As a goal directs processing, it produces storage side effects automatically that facilitate future goal achievement. As the goal directs processing, it also directs storage.

Failure-driven learning constitutes an interesting case of goal-driven storage. On the one hand, failures to solve a problem might not be stored, because these plans, in a narrow sense, are not goal relevant—technically speaking, they will not be necessary later for solving the problem. In a broader sense, though, failures are goal relevant, because they help avoid making the same mistake again (Schank, 1982). For this reason, a problem solver might decide to store failures. Note that these two options, as stated, assume implicitly that the problem solver has the ability to turn storage on and off, deciding when to store information based on its goal relevance. As we have seen, however, humans do not function in this manner, but simply store information as a side effect of processing.

Interestingly, the notion of storage side effects accounts naturally for failure-driven learning: Because all processing produces storage side effects, problem-solving failures, which are processing events, produce storage side effects, namely, memories of these failures. Without having to invoke any sort of decision-making apparatus for storage, we can account naturally for the availability of problem failures in memory that support failure avoidance. As this example illustrates, studying the nature of processing informs us about the nature of knowledge in memory and its origins.

17.4.2 Processing and Storage in Implicit Orientation

The processes of classic problem solving store information relevant to solving each particular type of goal individually; they do not store information relevant to achieving most other goals. Certainly, some transfer may occur from one goal to another, to the extent that their plans are similar, but in general storage is relatively goal specific, primarily facilitating the processing of the same goal on future occasions. Much work in the skill literature demonstrates the context-specificity of the knowledge stored during skill acquisition (e.g., Anderson, 1983; Ross, 1984; Shiffrin, 1988).

In contrast to classic problem solving, implicit orientation stores information that is more goal independent, information that can serve a much wider variety of goals. The basic idea behind implicit orientation is that an intelligent system develops and updates a world model (Barsalou, this volume; Barsalou et al., 1993; Leake, 1992). Most generally, a world model contains a person's beliefs about the current state of the world. More specifically, it contains a person's knowledge about places in the world and relations between them. For example, a person might believe that various countries, cities, neighborhoods, and houses exist, and also believe that they bear various relations to one another, such as adjacency, inclusion, and so forth. Embedded within this system of spatial frames are people's beliefs about the existence of specific objects and activities. For example, someone might believe that certain artifacts exist at certain locations in a house, and that certain activities are likely to be occurring in a particular office building.

How does implicit orientation bear on world models? One possibility is that the cognitive system has an implicit system goal of constantly orienting itself with respect to its world model. As people move about the world, they categorize locations, objects, and activities in the environment to orient themselves with respect to what they already know about the world. Certainly not everything in a current situation is identified during implicit orientation, given infinite entities could be identified. Depending on a variety of factors, such as perceptual salience and expectation, only a subset of the entities present are processed.

Based on those entities processed, the relevant part of a person's world model is retrieved to guide further processing. On recognizing a particular location, for example, information about it in the world model is retrieved, producing inferences about objects and activities likely to be present, as well as adjacent locations. Even though these locations, objects, and activities may be irrelevant to the strategic goal currently being pursued, information about them is nevertheless processed, because the system has a background goal of constantly orienting itself with respect to its world model.

How might the processing that underlies orientation inform us about the storage of information in a world model? Again, assume that processing stores a record of itself as a side effect. As each location, object, and activity in the world is categorized, a storage side effect ensues, storing a record of the identified entity in memory. In this way, implicit orientation processing stores

an updated record of the world as a side effect. Moreover, because only a subset of the entities in the current situation are processed, only records for this subset are stored in memory—not every possible entity present in the situation is updated.

In a given situation, it is likely that the relevant part of a world model typically becomes active and brought to bear on current orientation processing. For example, when a familiar location is recognized, information about it is retrieved from the world model and brought to bear on the categorization of current entities in it. Because this part of the world model is active as current entities in the world are being categorized, the storage side effects of these categorizations are integrated with the retrieved part of the world model. In this manner, updated information becomes stored with previous information about the same entities, simply because these new categorizations occur simultaneously with retrievals of old information from the world model (i.e., the one-entity one-frame principle in Barsalou et al., 1993). Once again, the structure of processing provides insights into the structure of knowledge.

Interestingly, the establishment of a world model in this manner is not a strategic goal. People usually do not have the intention to update world models and are not even aware that they are doing so. Moreover, much of the implicit orientation that takes place has nothing to do with the current strategic goals that people are pursuing. Instead, implicit orientation serves a system goal of maximizing the success of unanticipated goals that may arise later (Barsalou, this volume, chap. 5).

For example, imagine that I walk through a hotel lobby at a conference, having the strategic goal to exit the hotel. As I walk, I categorize a collection of chairs, sofas, and tables, as well as a convenience store and a bar. Although these entities and their locations to each other are irrelevant to my current goal, they are nevertheless stored as side effects of orienting myself within the environment. Several hours later, I may need a place where a colleague and I can discuss a research project. I can achieve this strategic goal easily if I can find a location in my world model that will enable it, otherwise I will have to search through the world physically to find one. Because implicit orientation caused relevant information about the lobby to be stored in my world model, it established a potential solution.

As this example illustrates, learning that served no strategic goal at the time later turns out to serve one. Moreover, this learning may support a wide variety of other strategic goals that are unrelated. For example, the presence of a seating area in the lobby could serve as a place to talk with an old friend, a source of tourist information, a place to relax and read the newspaper, a place to rendezvous with friends before dinner, and so forth. Unlike classic problem solving, which primarily stores information relevant to solving only a single type of problem, implicit orientation stores information relevant to solving a wide variety of unanticipated problems.

As we saw earlier for classic problem solving, the assess-search-execute cycle constitutes the basic form of processing, storing records of itself in

the process. For implicit orientation, the structure of knowledge is different, because the structure of processing is different. During implicit orientation, the cognitive system categorizes locations, objects, and activities into basic level categories (Barsalou, this volume, chap. 5). As many theorists have argued, such categorizations depend critically on analyses of physical structure (Biederman, 1987; Rosch et al., 1976; Tversky & Hemenway, 1985). In categorizing a chair, extracting the configuration of its seat, back, and legs is critical. Once an entity has been categorized on the basis of its physical structure, this structure is stored in a world model as a side effect. Most importantly, storing this structure enables functional inferences later that support a wide variety of unanticipated goals. From having stored the physical structure of a chair, one can infer numerous unanticipated functions, including its use as a doorstop, ladder, firewood, and so forth. Understanding the process of categorization yields insights into the information stored in world models as side effects.

17.5 CONCLUSION

As classic problem solving and implicit orientation illustrate, there is no single answer to the question of how goals constrain storage. Because each type of goal requires different processing, each type establishes different knowledge in memory as a side effect. This observation suggests one final point:

8. To the extent that each type of goal-driven behavior involves different cognitive processes, it will establish different knowledge in memory.

Rather than attempting to establish a single, uniform account of goal-driven storage, it will be necessary to develop many individual accounts. Each basic activity that constitutes intelligence must be examined individually to understand its unique processes, as well as its corresponding storage side effects. If we can understand how goals constrain processing, we put ourselves in a good position to understand how goals constrain learning. Once the diverse processes that underlie intelligence are understood, a comprehensive theory of learning will be significantly closer. Of course, other tools, such as a satisfactory theory of representation, will be necessary as well (Barsalou, 1993; Barsalou et al., 1993).

ACKNOWLEDGMENTS

This chapter was supported by funding from the Army Research Institute (MDA 903-90-K-0112).

REFERENCES

- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., & Bower, G. H. (1972). Recognition and retrieval processes in free recall. *Psychological Review*, 79, 97-123.

- Barsalou, L. W. (1992). *Cognitive psychology: An overview for cognitive scientists*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, & M. A. Conway (Eds.), *Theories of memories* (pp. 29–101). London: Lawrence Erlbaum Associates.
- Barsalou, L. W., Yeh, W., Luka, B. J., Olseth, K. L., Mix, K. S., & Wu, L. (1993). Concepts and meaning. In K. Beals, G. Cooke, D. Kathman, K. E. McCullough, S. Kita, & D. Testen, (Eds.), *Chicago Linguistics Society 29: Papers from the parasession on conceptual representations* (pp. 23–61). University of Chicago: Chicago Linguistics Society.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115–147.
- Bousfield, W. A. (1953). The occurrence of clustering in the recall of randomly arranged associates. *Journal of General Psychology*, 49, 229–240.
- Cofer, C. N. (1967). Does conceptual organization influence the amount retained in immediate free recall? In B. Kleinmuntz (Ed.), *Concepts and the structure of memory* (pp. 181–214). New York: John Wiley.
- Cohen, N. J., & Squire, L. R. (1980). Preserved learning and retention of pattern analyzing skill in amnesics: Dissociation of knowing how and knowing that. *Science*, 210, 207–210.
- 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.
- Eagle, M., & Leiter, E. (1964). Recall and recognition in intentional and incidental learning. *Journal of Experimental Psychology*, 68, 58–63.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1–67.
- Hammond, K. J. (1990). Case-based planning: A framework for planning from experience. *Cognitive Science*, 14, 385–444.
- Hyde, T. S., & Jenkins, J. J. (1969). Differential effects of incidental tasks on the organization of recall of a list of highly associated words. *Journal of Experimental Psychology*, 82, 472–481.
- Jacoby, L. L., & Brooks, L. R. (1984). Nonanalytic cognition: Memory, perception, and concept learning. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 18, pp. 1–47). New York: Academic.
- Jacoby, L. L., & Hayman, C. A. G. (1987). Specific visual transfer in word identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 456–463.
- Kolodner, J. L., & Simpson, R. L. (1989). The MEDIATOR: Analysis of an early case-based problem. *Cognitive Science*, 13, 507–550.
- Leake, D. B. (1992). *Evaluating Explanations: A Content Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive Psychology*, 12, 523–553.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492–527.
- MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 126–135.

- Mandler, G. (1967). Organization and memory. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 1, pp. 327–372). New York: Academic.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Puff, C. R. (1970). Role of clustering in free recall. *Journal of Experimental Psychology*, *86*, 384–386.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
- Ross, B. H. (1984). Reminders and their effects in learning a cognitive task. *Cognitive Psychology*, *16*, 371–416.
- Scarborough, D. L., Cortese, C., & Scarborough, H. S. (1977). Frequency and repetition effects in lexical memory. *Journal of Experimental Psychology: Human Perception and Performance*, *3*, 1–17.
- Schachter, D. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 501–518.
- Schank, R. C. (1982). *Dynamic memory: A theory of reminding and learning in computers and people*. New York: Cambridge University Press.
- Shiffrin, R. M. (1988). Attention. In R. C. Atkinson, R. J. Herrnstein, G. Lindzey, & R. D. Luce (Eds.), *Stevens' handbook of experimental psychology: Vol. 2. Learning and cognition* (pp. 739–811). New York: John Wiley.
- Squire, L. R. (1987). *Memory and brain*. New York: Oxford University Press.
- Stroop, J. R. (1935). Studies on interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662.
- Tulving, E. (1962). Subjective organization in free recall of "unrelated" words. *Psychological Review*, *69*, 344–354.
- Tulving, E. (1964). Intratrial and intertrial retention: Notes toward a theory of free recall verbal learning. *Psychological Review*, *71*, 219–237.
- Tulving, E. (1966). Subjective organization and the effects of repetition in multi-trial free recall verbal learning. *Journal of Verbal Learning and Verbal Behavior*, *5*, 193–197.
- Tversky, B., & Hemenway, K. (1985). Objects, parts, and categories. *Journal of Experimental Psychology: General*, *113*, 169–193.