THE UPDATING OF HUMAN MEMORY

Robert A. Bjork[1]

UNIVERSITY OF CALIFORNIA, LOS ANGELES

I. Introduction

Watson: Perhaps you will kindly give me a sketch of the course of events from memory.

Sherlock Holmes: Certainly, though I cannot guarantee that I carry all the facts in my mind. Intense mental concentration has a curious way of blotting out what has passed. The barrister who has his case at his fingers' ends and is able to argue with an expert upon his own subject finds that a week or two of the courts will drive it all out of his head once more. So each of my cases displaces the last, and Mlle. Carere has blurred my recollection of Baskerville Hall. Tomorrow some other little problem may be submitted to my notice which will in turn dispossess the fair French lady and the infamous Upwood. So far as the case of the hound goes, however, I will give you the course of events as nearly as I can and you will suggest anything which I may have forgotten.

A. Conan Doyle, *The Hound of the Baskervilles*

[1]This chapter was completed while the author was a guest of the Human Information-Processing Department, Bell Laboratories, Murray Hill, New Jersey. The research reported in Section III,B,2 was supported by a grant from the Research Committee of the University of California, Los Angeles.

ISBN 0-12-543312-3

A. GENERALITY AND IMPORTANCE OF UPDATING

Everyday functioning requires that we keep our memories reasonably current. To the degree that we do not somehow set aside or eliminate information no longer needed, we become confused, error-prone, and inefficient. The following brief story, though of questionable literary merit, illustrates some typical consequences of failing to update efficiently.

> The day got off to a bad start for Professor Sterling Theorist when he was late for the first meeting of his Statistics 250B course in the Winter Term. He had forgotten that the Elm Street Exit was temporarily closed for repairs, which cost him about 10 min in backtracking. Even then, he would not have been excessively late had he not first gone to the room in which he had taught Statistics 250A in the fall. After struggling through his lecture, Sterling retreated to the Department lounge to recover over a cup of coffee, where he found Professors Wil Parish and Grant Funding discussing the 1977 World Series. He interrupted their conversation to tell them a story about George Steinbrenner, owner of the New York Yankees, who supposedly was outside Yankee Stadium when a vendor asked him if he wanted to buy a pennant. Steinbrenner replied, "No thanks, I already bought one."
>
> "That's a good joke, Sterling," responded Professor Parish, "but I liked it better the first time you told it to us."
>
> In order to change the subject, Theorist asked Professor Funding, Chairman of the Department, if he knew who was scheduled to be the first colloquium speaker in the winter term. Professor Funding replied that he wasn't sure, to which Theorist remarked, "Well, I sure hope he's more comprehensible than that last turkey we had in the fall."
>
> "*I* was the last turkey to speak in the fall," said Professor Funding.
>
> After an awkward effort to explain that he was referring to another speaker, one who he thought had been the last to speak in the fall, Sterling retired to his office. He managed to avoid further problems until his wife called at 5:45 P.M. to tell him he was supposed to be home 15 min ago because they had agreed to celebrate the 6-month anniversary of their wedding. "Oh, no," said Theorist, "I'll be right home!"—which he was not, however, because it took him an extra 10 min to find his car in the strange part of the parking structure to which he had been relegated that morning owing to his late arrival. When he finally arrived home, Theorist found a very angry wife in tears sitting at the kitchen table with a very flat souffle sitting on the counter.
>
> His wife looked at him and said, "Sterling, how could you?" whereupon Sterling replied, "But Sheila—I mean Shirley!" At that point, however, having made the monumental blunder of referring to Shirley by his first wife's name, explanations were no longer possible, and nothing remained to be salvaged from Sterling's miserable day.

In the foregoing brief vignette, Professor Theorist suffers no less than seven different updating failures. Most of us, happily, do not have such days very often, if ever; nor are we very likely to commit such apocalyptic updating errors as referring to one's current spouse by the first name of one's former spouse. In general, however, the kind of updating problems that bedevil Professor Theorist are more familiar than fictional to all of

us. We need to remember where we left the car *today*, we need to remember our *current* phone number, and we need to remember what the trump suit is on *this* hand. In adding a column of numbers on a calculator, we need to be clear about what number is current, and so forth. Thus, in any number of situations, on any number of time scales, we need to discriminate current to-be-remembered information from out-of-date, to-be-forgotten information.

The generality and importance of updating are also apparent when one considers that updating processes are intrinsic to job environments ranging from short-order cook to intelligence analyst. An air-traffic controller, for example, typifies one type of job where updating is crucial. At any one point in time, an air-traffic controller is responsible for a set of information that denotes the status of some number of planes. At some later point that set will be replaced by a new set of information, and it is highly desirable that the controller not be confused as to the set membership of any given item of information. The pilot's task as well requires continual updating of heading, altitude, speed, and so forth. People in command and management positions also bear formidable updating burdens, particularly in crisis situations, as do individuals whose job it is to keep track of the present status of supplies, parts, personnel, and equipment. Finally, any change in hardware or software requires an updating of skills or procedures. In all of these cases, of course, we entrust much of the updating burden to external memories, such as computers or pencil and paper. Even the best of cooks, controllers, and commanders, however, remain susceptible to updating errors owing to momentary information overloads and the frailties of the human memory system.

B. ASPECTS OF THE UPDATING PROBLEM

This chapter focuses on certain selected aspects of the updating problem. In the Section II, I shall discuss the roles of encoding processes and contextual factors in updating. Section III deals with two aspects of the breakdown of updating: the reinstatement of to-be-forgotten information and the regression of memory for to-be-remembered information. Finally, in Section IV, I shall attempt to look at the updating problem from the standpoint of interference theory.

II. How Do We Update?

The need to update poses a fundamental memory problem. On the one hand, as pointed out above, it is helpful to forget or set aside information that is no longer current. On the other hand, we may later want to retrieve out-of-date information, as in Sherlock Holmes's attempt to recall the details of the Baskerville case, or we may need to remember past information in order to interpret current information properly. Different updating processes clearly differ in how much the act of updating destroys the past. Computers, for example, employ a quite radical updating mechanism: when new information is stored at a memory location, the old information at that location is obliterated. The displacement mechanism that has been offered to explain forgetting from short-term memory is another example of an updating process that destroys the past. Whatever the mechanisms involved in human long-term memory, they are clearly not so destructive (nor so efficient). We can remember who won the last Super Bowl football game without forgetting who won all the preceding Super Bowl games (although we may be confused, at least momentarily, as to which game was the last game). Sherlock Holmes advances the theory that intense mental concentration on current information will obscure or "blot out" past information. He then goes on to show, however, that he in fact retains many of the details of the no-longer-current Baskerville case. Section II,A reports an effort by R. A. Bjork and McClure (1974) to examine experimentally the consequences of several different updating strategies.

A. ROLE OF ENCODING PROCESSES

Bjork and McClure distinguish between *destructive updating* and *structural updating*. Structural updating, in contrast to destructive updating, does not involve the destruction of past information. Rather, successive inputs are encoded as a series in which some underlying structure specifies which input is most recent. Thus, one's ability to give back the last word of a spoken sentence is in part determined by syntax or phrase structure. The more there is some principle that connects or orders successive inputs, the fewer are the chances that order information will be lost. If there is little or no superordinate structure, however, order information is lost rapidly. Thus, at any point in time, it would be difficult for a short-order cook to reconstruct the series of orders he has completed.

The updating task employed by R. A. Bjork and McClure (1974) was a continuous paired-associate task in which a series of response words was associated with each of four different stimulus words. At any point in the task, subjects were responsible for remembering the last response word paired with each stimulus word. The actual experiment consisted of a series of trials, each of which had the following structure: one of four stimulus words was presented, the subject attempted to recall the last response word paired with that stimulus, and then a new to-be-remembered response word was presented together with the stimulus word.

This paradigm goes back at least as far as a study by Yntema and Mueser (1960) and is sometimes referred to as the maximal PI (proactive interference) paradigm because retrieval of the current response associated with a given stimulus is susceptible to PI from all prior responses associated with that stimulus. In fact, Atkinson and Shiffrin (1968) used the maximal PI paradigm in an attempt to obtain a short-term retention function that would be uncontaminated by long-term memory. They assumed that the repeated association of new response words to a given stimulus would render long-term memory essentially useless. Whether they were right or wrong in that assumption depends, according to Bjork and McClure's characterization, on the nature of subjects' encoding processes. Destructive updating would, of course, render long-term memory useless, whereas structural updating would clearly not. A nondestructive updating process that was also nonstructural would not yield very useful information in long-term memory, unless differences in some strength measure or temporal tagging could be utilized to infer something about input order.

The basic procedure in Bjork and McClure's study was quite straightforward. Subjects were given three decks of cards, a practice deck and two experimental decks. They went through each deck turning over the cards one at a time. Each deck consisted of a series of test-study trials. The test and study phases of a given trial were on separate cards as shown in Fig. 1. After a particular stimulus and response were presented together for study, 0, 1, 2, 6, or 10 test-study trials involving other stimuli intervened before that stimulus was presented again as a probe of the subject's memory for the response member of the pair. Thus, in Fig. 1, the stimulus word FROG is paired with WALL, then PLUM, and then COAT in the sample segment shown.

At the end of the experiment, without forewarning, subjects were handed a sheet of paper and were asked to write down all the response words they could remember. The sheet was divided into five columns, the first four of which were headed by the four stimulus words (BOAT, ROPE, HILL, and FROG), and the last of which was headed by a row of question marks. Subjects were asked to write a given response word in the column headed by the stimulus with which it had been paired; if they were unable to remember which stimulus word had been paired with that response word, they were asked to write the response word in the column headed by question marks.

Each subject was asked to use one of three different encoding strategies as he or she went through the decks. One of these strategies was designed to yield essentially destructive updating, another was designed to yield structural updating, and the third was designed to fall somewhere in between.

As a destructive-updating strategy, Bjork and McClure first attempted

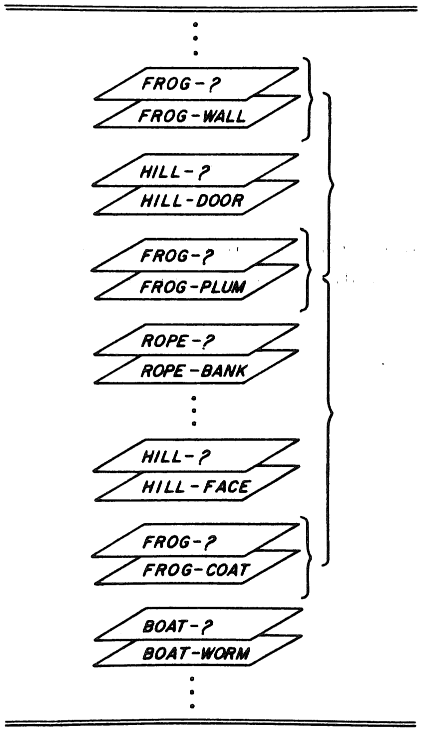

Fig. 1. Sample sequence of trials in Bjork and McClure's (1973) updating experiment. (From R. A. Bjork & McClure, 1973.)

to have subjects use a strategy supposedly used by V.P., a man whose extraordinary memory abilities were studied in considerable detail by Hunt and Love (1972). One of the array of tasks administered to V.P. was the maximal PI task. When V.P. performed without error on an initial version of the task, Hunt and Love made the task more difficult, but V.P. was still essentially perfect. When asked how he approached the task, V.P. said that he used a visual-erasure strategy (Hunt, personal communication). V.P. said that he imagined that there were several little blackboards on the wall in front of him, each with the name of a stimulus above it. When he studied a response to a particular stimulus, he first imagined erasing whatever was on the blackboard, and he then imagined writing the new response on the blackboard. When he was tested with a particular stimulus, he simply "read" the response written on the imaginary blackboard corresponding to that stimulus. Assuming that V. P. was doing what he said he was doing, his visual-erasure strategy is a near-perfect example of destructive updating. Had he been asked, at the end of

the experiment, to recall as many responses as he could, his performance should have been dismal—which might have been quite a shock for him.

Bjork and McClure were completely unsuccessful in having the garden-variety subject use the visual-erasure strategy. For all subjects, whatever the strategy they were asked to use, the four stimulus words (BOAT, ROPE, HILL, FROG) were printed on 5 × 7 cards that were mounted in a horizontal array on a wall directly in front of the subject. Even with that aid, however, subjects (with one exception) could not begin to do the imaginal writing and erasing. The one exceptional subject that Bjork and McClure encountered, who could carry out the visual-erasure strategy, though less than perfectly, may indicate that there are other people in the world besides V.P. who have the visual-imagery capacity to employ his system.

As an alternative to the visual-erasure strategy, Bjork and McClure used an ordered-rehearsal strategy. In the ordered-rehearsal condition, subjects were asked to rehearse the four response words they were responsible for at any point in the task in a fixed, rote order corresponding to the order of the stimulus words on the wall. Thus, when one of the stimuli was presented as a probe test, the subject would give the response word corresponding to that stimulus and when the new response word was presented, the subject would insert the new word in the correct place in his rehearsal scheme. The ordered-rehearsal condition was deemed to approximate destructive updating for two reasons: (*a*) new items displace old items in a short-term rehearsal set, and (*b*) rote, ordered rehearsal has been shown to have little or no effect on long-term recall (see, e.g., Craik & Watkins, 1973; Woodward, Bjork, & Jongeward, 1973).

As a structural-updating strategy, a simple story-construction strategy was employed. The subjects in the story-construction condition were asked to construct a continuing narrative based on each stimulus word. Thus, in the case of the stimulus word FROG, a subject might use the successive response words paired with FROG to compose an adventure story involving a frog.

As a strategy that was in between in the sense of not being obviously destructive or structural, an image-replacement strategy was employed. When a response word was presented for study with a particular stimulus, the subjects were to imagine engaging in a particular activity. In the case of ROPE, for example, they were asked to imagine tying rope around whatever was denoted by the response word. When the next response word was paired with rope, the subject was asked to imagine untying the rope from whatever it was tied around, and to imagine tying it around whatever was denoted by the new response. For FROG, BOAT, and HILL as well, the subjects were asked to employ specific imaginal activities. In all cases, there was a definite replacement required, that is, an undoing followed by a doing. The image-replacement strategy was regarded as "in between" because it was not obvious whether such a strategy would amount to destructive updating, or rather, would result in a series of images (possibly poorly ordered) in memory.

The basic results of Bjork and McClure's experiment are shown in Fig. 2. The story-construction strategy was effective both in terms of the within-deck updating and in terms of the final total recall. Apparently, the story narrative corresponding to each stimulus word functioned as a structure that not only facilitated retrieval from long-term memory, but also provided a basis for judging which response word was most recent.

The most interesting aspect of the results in Fig. 2 is the disordinal interaction in recall performance for the ordered-rehearsal and image-replacement conditions as a function of time of test. The ordered-rehearsal condition resulted in better updating performance and substantially poorer final recall than did the image-replacement condition. Thus, the ordered-rehearsal condition was quite effective in eliminating PI from prior response words, but it largely eliminated those words from final recall as well. (The fact, however, that over 20% of the response words were still available at the time of the final recall test demonstrates that the ordered-rehearsal strategy was not a completely destructive updating process.) The image-replacement strategy, on the other hand, appeared to

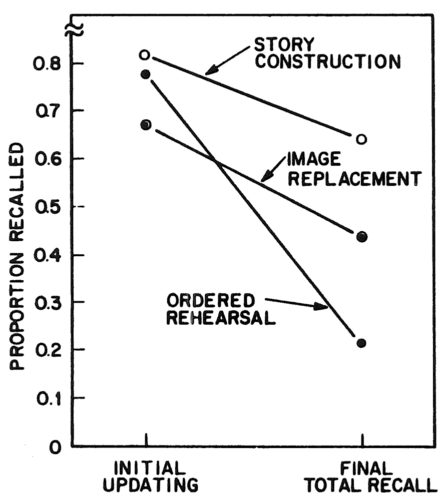

Fig. 2. Proportion correct responses on tests of initial updating and final recall as a function of encoding strategy. (From R. A. Bjork & McClure, 1973.)

yield a more accessible series of images in memory, but the ordering of those images was quite fallible.

The final-recall results shown in Fig. 2 are based on free-recall scoring. A word was counted as correctly recalled if it appeared anywhere on the final-recall sheet. Had response words been counted correct only if they appeared in the column headed by the stimulus with which they had been paired, the advantage of the story-construction and image-replacement conditions over the ordered-rehearsal condition would have been magnified. In the story-construction and image-replacement conditions, response words were almost always recalled in the correct column, whereas response words in the ordered-rehearsal condition were frequently written in the wrong column. (The conditional probabilities that a word recalled on the final test was recalled in the correct column were .97, .95, and .71 for the story-construction, image-replacement, and ordered-rehearsal conditions, respectively.)

One additional analysis of the final-recall data from the Bjork and McClure study merits comment. Bjork and McClure computed the rank-order correlation between the order in which recalled response words were presented during input and the order in which the were recalled during final recall. That analysis produced an interesting result: Whereas the correlation averaged across individual subjects in the ordered-rehearsal and image-replacement conditions was distinctly negative (−.437 and −.417, respectively), the average correlation in the story-construction condition was distinctly positive (+.513). Thus, in both the ordered-rehearsal and image-replacement conditions, it appears that the long-term memory representation resulting from those strategies was most efficiently searched by working backward in time. The story-construction strategy, on the other hand, appears, reasonably enough, to have resulted in a memory representation from which it was natural to retrieve words working forward in time.

One implication of Bjork and McClure's research is that the "goodness" of any particular encoding strategy is not an absolute matter. Thus, on the basis of total long-term recall, one might judge the image-replacement strategy to be better than the ordered-rehearsal strategy, but if one's concern was to optimize ongoing performance rather than long-term recall, as it might be in a number of job settings, the ordered-rehearsal strategy is clearly the better of the two. Similarly, Bjork and Jongeward (reported in R. A. Bjork, 1975) have shown that as a short-term holding operation, rote (primary) rehearsal is superior to elaborative (secondary) rehearsal. Thus, if one's only goal with respect to a phone number, for example, is to maintain that number in memory long enough to walk across the room and dial it correctly on the phone, rote, cyclic rehearsal is better advised than is an effort to note patterns in the number, similarities to other numbers in one's long-term memory, and so forth (even though such mnemonic activities may greatly improve one's long-term recall of the number).

B. ROLE OF RETRIEVAL CONTEXT

How well current to-be-remembered information is retrieved depends not only on the type of encoding used at the time of storage, but also on how well the cues at retrieval match the cues that were present at storage (see Tulving and Thompson's 1973, "Encoding Specificity Principle"). Many cases where out-of-date, to-be-forgotten information is retrieved in place of current information are traceable to the fact that the local retrieval context matches the storage context of the past information better than it does the storage context of the current information.

To illustrate the latter point, consider the following anecdote. A professional couple, before they had children, used to work at their office routinely until 1 or 2 A.M., at which point they would drive home to their apartment. They then had a child and bought a house on the other side of town. Because their child had the typical infant's preference for arising early in the morning, their work schedules as well as their home address changed. After a full year of the new regimen, the husband was forced one evening to work at his office until 2 A.M. He then left the deserted building and drove home, but the behavior pattern he retrieved from memory led him almost to the door of his old apartment before he realized that something was not quite right. Most of the situational cues that were available when he left his office were ones that were associated with the out-of-date behavior pattern rather than with the appropriate behavior pattern.

Smith, Glenberg, and Bjork (1978) have recently reported a series of experiments in which they examine the influence of situational context on recall and recognition. One of their experiments demonstrates quite nicely that both context and recency are important factors in determining what items are accessible in memory.

Smith *et al.* had subjects study a first paired-associate list on Day 1 in a particular physical context. On Day 2 the subjects studied a different paired-associate list in a different physical context. Each paired-associate list consisted of 45 word–word pairs in which the stimulus word was a weak associate of the response word (e.g., car–BODY). The two lists overlapped in only one respect: 15 of the stimulus words were common to both lists. Each list was presented four times in random order, and at the end of each session 15 stimuli that were unique to the list studied in that

session were presented as a probe test of subjects' memory for the associated responses. The purpose of this partial test was to give some sense of closure at the end of each study session.

On the third day subjects were brought back and tested in either the Day 1 context, the Day 2 context, or a neutral context. The 15 common stimuli were presented together with the remaining 15 unique stimuli from each list. The test was carried out as an MMFR test, that is, next to each stimulus subjects were asked to write as many responses as they could remember having been paired with that stimulus.

The Day 1 and Day 2 contexts were designed to differ from each other in a number of ways. One of the two contexts was a small, windowless room in an old off-campus building. There was a large blackboard, several cabinets, and a high level of general clutter. The experimenter was dressed in a coat and tie, and the paired associates were shown on slides. The other context was a tiny room within the animal laboratories of a large, modern, central-campus building. Two windows overlooked a courtyard, and a one-way mirror covered one wall. The experimenter was dressed in a flannel shirt and jeans, and the paired associates were presented via a tape recorder. The neutral context was a large classroom overlooking a busy street. The experimenter wore routine, nondescript clothes, and the stimuli were read aloud and simultaneously shown on index cards.

The results of Smith *et al.*'s experiment are shown in Fig. 3. In the top panel the proportion of cases in which the Day 1 or Day 2 responses were recalled in reaction to the unique stimuli is shown as a function of the test context on Day 3. The bottom panel shows the same results for the common cues.

There are three aspects of the results in Fig. 3 that merit comment. (*a*) There is a recency effect; when the Day 3 test took place in a neutral context, Day 2 responses were more frequently recalled than were Day 1 responses. The human memory system does not work like a simple push-down stack, however, because (*b*) there is also a context effect. The interaction of the solid lines in both panels demonstrates that responses whose storage and retrieval contexts matched were recalled more frequently than responses whose storage and retrieval contexts mismatched. (*c*) Finally, there is something of an anomaly in the recall of Day 1 responses. The three points on the right-hand side of each panel (i.e., the proportions of Day 2 responses that were recalled) are ordered the way one would expect: Recall is best when the test took place in the Day 2 context, worst when the test took place in the Day 1 context, and in between when the test took place in the neutral context. The left-hand points in both panels, however, reveal that Day 1 responses were recalled

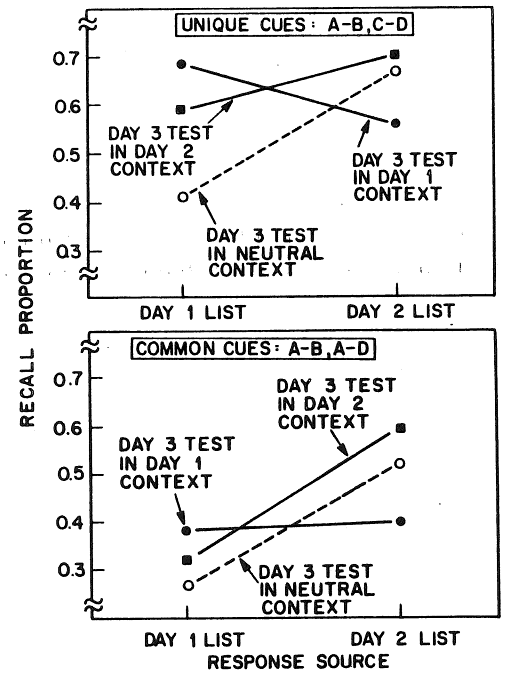

Fig. 3. Proportion of List 1 and List 2 responses recalled correctly as a function of Day 3 test context. (From Smith, Glenberg, & Bjork, 1978).

less frequently when the test took place in the neutral context than they were when the test took place in the Day 2 context. A possible explanation of this apparent anomaly is that subjects, because they had no particular reason not to, reinstated (retrieved) some of the Day 1 responses during the Day 2 study session. Such responses then became associated to the Day 2 context as well as the Day 1 context, which would possibly explain why the recall of Day 1 responses is better in the Day 2 context than it is in the neutral context. Such reinstatement processes are one of the topics I shall address in the next section.

III. Breakdown of Updating

Updating, even though apparently well-established, may fail or break down at some later point under certain circumstances. One such circumstance was illustrated above by the anecdote about the husband who

mistakenly drove home to an apartment he had not lived in for a year. That anecdote demonstrates that when the local retrieval context matches the storage context of out-of-date information better than it does the storage context of currect information, we become susceptible to intrusions of out-of-date information. In this section I shall discuss two other processes that may be responsible for the breakdown of updating: the reinstatement of to-be-forgotten information and regression of memory for to-be-remembered information.

It is worth commenting that the breakdown of updating can be a formidable problem from a practical standpoint. Even when people are apparently well-trained in the use of new equipment or procedures, intrusions of behaviors appropriate to the old equipment or procedures still occur, often at costly times. Such intrusions are particularly likely to occur in emergency situations. A person may have owned his new car for a year, for example, but when the brakes fail he finds himself reaching for the emergency brake in the place it was in his old car rather than in its current place. The military has been troubled by this kind of problem for years, with pilots under stress making inappropriate responses in new aircraft, and so forth. It may be that emergency situations constitute a kind of context that is more associated with the old equipment or procedures than it is with the new equipment or procedures. It is not easy to incorporate realistic emergency situations in training programs.

A. REINSTATEMENT OF TO-BE-FORGOTTEN INFORMATION

The passage from *The Hound of the Baskervilles* quoted at the outset of this chapter is preceded by a passage in which Dr. Watson says,

> My friend was in excellent spirits over the success which had attended a succession of difficult and important cases, so that I was able to induce him to discuss the details of the Baskerville mystery. I had waited patiently for the opportunity, for I was aware that he would never permit cases to overlap, and that his clear and logical mind would not be drawn from its present work to dwell upon memories of the past.

Sherlock Holmes clearly holds the view that in order to optimize the processing of current information, one should not resurrect or reexpose oneself to potentially interfering past information. That Sherlock Holmes has a point is demonstrated by the results of a study by E. L. Bjork, Bjork, and Glenberg (1973).

In Bjork *et al.*'s experiment, subjects were presented with three types of lists. In one type of list (FR lists), subjects were shown 16 common words one at a time, followed by a signal (a row of minus signs) to forget those words. They were then shown a second set of 16 words one at a time. A second type of list (R′R lists) was like the FR lists except that the signal (a row of plus marks) instructed subjects to remember rather than to forget the first 16 words. The final type of list [(−)R lists] was also like an FR list except that there was no first set of 16 words. Rather, there were 16 slides each of which showed a pair of figures that differed from each other on zero, one, two, or three stimulus dimensions (e.g., size). For each slide, subjects simply had to say how many stimulus attributes were common to the two figures. After the 16 figure slides, a row of minus signs appeared as in the FR lists, and then a set of 16 words was presented. The figure task was included in the design to insure that the general sequence and timing of events in the (−)R condition was the same as the sequence of events in the FR and R′R conditions. Subjects were presented nine lists, three of each type. The ordering was haphazard so subjects did not know in advance what type of list was upcoming.

Subjects' recall of each type of list was tested in three different ways. A recall test was administered either immediately, after a forced-choice recognition test was completed, or after an arithmetic task was completed. In the FR and (−)R conditions, subjects were asked to recall the last 16 words in the list [that is, the only words presented in the (−)R lists and the words presented after the forget instruction in the FR lists]. In the R′R condition subjects were asked to recall all 32 words from the list.

The forced-choice recognition test consisted of eight pairs of words. In each pair, subjects were asked to circle the word they thought had been presented in the second sublist of 16 words. For FR and R′R lists, four of the distractor words were from the first sublist of 16 words and four were new words that had not appeared elsewhere in the experiment. For (−)R lists, all eight distractor words were new words. The arithmetic task was designed to take about the same length of time it took the typical subject to complete the recognition test.

The results of principal present interest are shown in Fig. 4. The three panels correspond to the three recall conditions, and in each panel the proportion of words recalled from the second 16-word sublist is plotted as a function of list type. In the right-hand panel, the proportions plotted are for only the eight words in the second sublist that *did not appear on the recognition test*.

The three panels in Fig. 4 appear to tell an interesting tale. In immediate recall (the left-hand panel), the forget instruction essentially eliminated PI owing to the first sublist. Performance in the R′R condition was significantly worse than performance in either the FR or (−)R conditions, which in turn did not differ significantly from each other. When recall was delayed by an arithmetic task (middle panel), there was an overall decrease in recall but no recovery of interference owing to the first

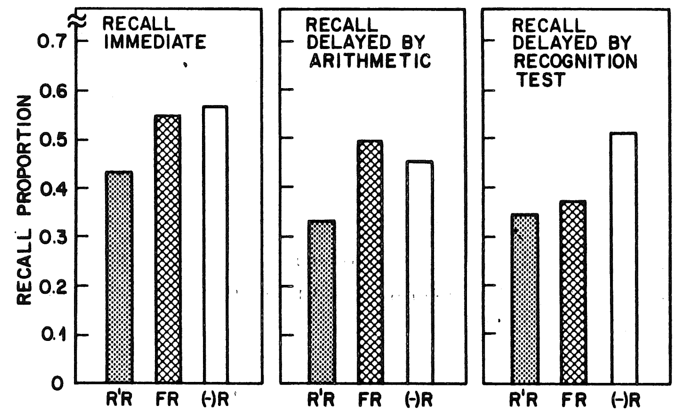

Fig. 4. Recall performance on the second sublist of 16 words as a function of list type and test condition; FR denotes lists in which the first sublist of 16 words was to be forgotten, R′R denotes lists in which the first sublist was to be remembered, and (−)R denotes lists where there was no first sublist of words. (After E. L. Bjork, Bjork, & Glenberg, 1973.)

sublist in the FR condition. Once again, recall in the R′R condition was significantly worse than recall in either the FR or (−)R conditions, which do not differ significantly from each other. The right-hand panel, however, presents a different picture. When recall was delayed by the forced-choice recognition test, there was a recovery of interference owing to the first sublist in the FR condition, even for those words that did not appear on the recognition test. Recall in both the R′R and FR conditions was significantly worse than recall in the (−)R condition, and the R′R and FR conditions did not differ significantly from each other.

Bjork *et al.* interpret the results in Fig. 4 as demonstrating that even the minimal exposure of four to-be-forgotten items as distractors on the recognition test can reinstate the interference attributable to the to-be-forgotten set. It was also the case that intrusion of to-be-forgotten words increased following the recognition test, even if one exludes from the analysis the four to-be-forgotten words that appeared on the recognition test. Thus, there appears to be a reinstatement of the entire to-be-forgotten set. Sherlock Holmes was apparently well-advised not to let his mind "be drawn from its present work to dwell upon memories of the past." This general topic arises again in the last section of this chapter.

B. REGRESSION OF HUMAN MEMORY

This section argues for what I believe is an absolutely fundamental property of human memory: that one's memory representation of a past experience, person, or situation regresses over time. To be more specific, assume that one has an ongoing relationship with another individual during a given period of one's life, and that at some point, for whatever reason, that relationship is discontinued. Immediately following the end of that relationship, one's memory image of that individual is clearly dominated by the most recent version of that individual. If, for example, the individual in question was a high school friend one had known since grade school, one's memory image would be dominated by the high school graduation version of that individual. With the passing of time, however, I think one's memory image regresses toward an earlier or more average version of that individual. Similarly, one's image of a town or a school where one had spent some time would regress toward an earlier or more average version of that town or school as the length of time one was away from that town or school increased.

It is a fairly common experience that one is surprised how much a child has grown up, a friend has aged, or a town has changed since the last time one saw that child, friend, or town. Such surprises might be largely interpretable in terms of the regression of memory representations. Children do grow up, of course, friends do age, and towns do change, but a subjective judgment of such changes based on the difference between a regressed memory representation and the current state of the child, friend, or town will overestimate the actual changes. One particularly compelling instance of such overestimation, in my view, occurs when one is away from one's small children for a week or two. The apparent growth can far exceed any actual growth that could have taken place in that time. The fact that a day or so later such phenomenal growth is no longer apparent argues against the reality of such apparent changes.

1. Possible Forms of the Regression Process

If the regression of human memory is as fundamental as I believe it is, then any memory representation that is not actually in the process of being updated, used, or rehearsed is, in fact, becoming less current. What might be the form of such a regression process? Two possibilities are the following. First, if we assume that a changing person or situation has left a series of representations in one's memory, it might be that the more recent representations differentially weaken with the lapse of time, or that earlier representations spontaneously recover in strength with time, or both. The net effect would be that the dominance pattern would change gradually with time: one's dominant memory representation of a person or situation would gradually move toward an earlier version of that person or situation. Such a mechanism would be not unlike the interference-theory explanation of PI in terms of unlearning plus spontaneous recovery. At the time of input, current representations suppress or inhibit earlier representations. With the passage of time, more recent representations weaken and earlier representations recover in strength.

A second possibility is that one's memory does not regress toward a specific earlier representation of a person or situation but, rather, toward a more average or prototypical version of that person or situation. Assume that during the course of one's experience with a changing person or situation two kinds of representations are left in memory. Each exposure leaves a specific, detailed representation in memory (at least temporarily), and from the series of exposures is abstracted a conceptual or prototypical representation of that person or situation. Assume further that the later type of memory representation is more durable or less volatile, that is, less subject to loss from memory than is the former type of representation. With the lapse of time, then, the specific, detailed representations of recent exposures are lost and one's memory regresses toward the abstracted average representation, which is in some sense, of course, ''younger'' than the most recent actual version of what it represents.

2. A Preliminary Study

Together with Douglas Nies at UCLA, I have initiated research on regression processes in face recognition. The initial experiment was quite simple. On each of a series of trials in the experiment, subjects went through the sequence of events illustrated in Fig. 5. After a READY signal, eight facial photographs of a particular boy or girl were shown in succession at a 2-sec rate. Across the eight photographs, the boy or girl's face aged from about the first-grade level to about the twelfth-grade level. After the last picture, a row of asterisks was presented for either 3 sec (short delay) or 12 sec (long delay). Finally, a test face was presented. The test face was either one of the eight photographs shown on that trial (in which case subjects were to press the YES button), or it was a photograph not presented elsewhere in the experiment of a generally similar boy or girl (in which case subjects were to press the NO button).

The idea behind the experiment was quite straightforward. We wanted to see if we could demonstrate a regression-type phenomenon in this miniaturized simulation of a real-world situation where we think memory regression does take place. The primary data of interest are subjects' reaction times. When tested at the short delay, we might expect that the subjects' memory representation would be dominated by the most recent (and oldest) of the pictures. The YES reaction times should then show a recency effect; that is, subjects should be faster at recognizing later pic-

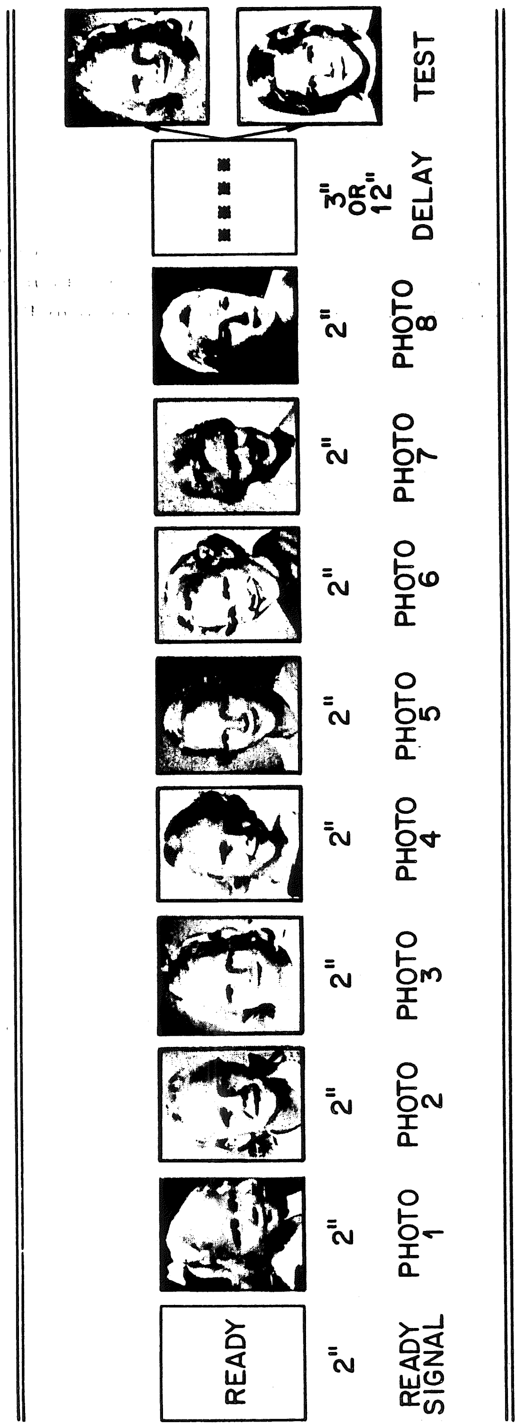

Fig. 5. Sequence of events on a typical trial.

tures in the series than they are at recognizing earlier pictures in the series. But what should happen when subjects are tested after the long delay? If during the delay there is a regression of subjects' memory for the boy or girl shown on that trial, we should no longer expect a recency effect. If such regression took the form of a regression to an earlier specific photograph in the series, we might actually find that the minimum reaction time was at some point earlier in the series than the most recent (oldest) photograph. On the other hand, if the regression was toward an average or prototypical representation of the boy or girl shown on that trial, it is a bit more difficult to say what we should expect. If the central person-concept abstracted from the series were to correspond more closely to the intermediate photographs than to the earliest or latest photographs, then we still might expect the minimum reaction time to fall at some intermediate point. At the very least, the recency effect should flatten out at the long delay.

The obtained reaction times for correct YES and NO responses are shown in Fig. 6 as a function of input serial position and delay of the test probe. The overall error rate was 8%. The 96 subjects in the experiment generated 48 observations per point in Fig. 6. Because the data were quite

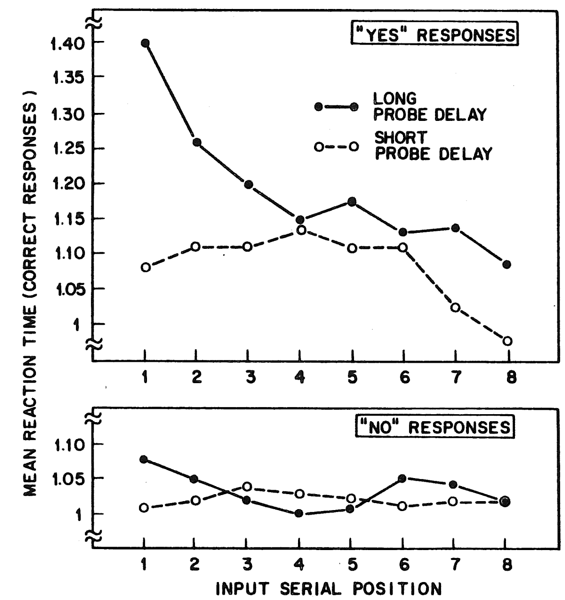

Fig. 6. Average reaction times for correct YES and NO responses as a function of probe delay and probe age.

noisy, the actual curves shown in Fig. 6 are "smoothed" curves, where every point except those at Position 1 and Position 8 is plotted as the average of three successive points.

The data in Fig. 6, though preliminary and probably not the best basis on which to speculate, are nonetheless suggestive. First of all, the YES reaction times at the short delay exhibit the recency effect one would expect. At the long delay, there is certainly no new minimum at an intermediate point in the YES reaction times, but there is, at least in the author's eyes, a regression effect of sorts. If one looks at the differences between the two YES curves in the top panel, then it is apparent that the intermediate pictures were recognized almost as rapidly at the long delay as they were at the short delay, whereas recognition of the most recent pictures was slowed down somewhat more and recognition of the earliest pictures was slowed down a great deal. The very slow YES reaction times for the youngest faces at the long delay might make sense in terms of the regression-to-a-prototype idea. As one looks through the pool of photographs used in the experiment, the youngest faces seem less differentiated than do the older faces. The features are less pronounced and there are fewer distinguishing characteristics, such as an unusual hair style. If the youngest face or two in a series, therefore, differ more from the abstracted prototype than do the later faces, the increase in the time taken to recognize those faces at the long delay is consistent with the idea that memory regression takes the form of regression to a prototype.

The NO reaction times in the bottom panel of Fig. 6 were faster than the YES reaction times and did not vary systematically as a function of either the delay or age of the probe face. It is as though the NO responses, whatever the age or delay of the test face, were always based on a categorical judgment that the test face differed from the prototype abstracted from the eight faces shown on that trial.

The preceding experiment on memory regression was included in this chapter because of its pertinence and stimulation value. Once again, the preliminary nature of the experiment should be emphasized. The astute reader has no doubt noticed that certain interesting control conditions were missing from the design and, as mentioned above, the data are quite noisy.

IV. Updating and Interference Theory

The problems of updating human memory can clearly be viewed as problems of interference and transfer. It is, therefore, not surprising that

the issues that arise in the analysis of updating are closely related to the issues that have been central to the analysis of PI over the last several decades. In this section, I shall attempt to relate some of the theoretical constructs that comprise interference theory to the aspects of updating outlined above.

A. ENCODING, UNLEARNING, AND CONTEXT

The distinction between destructive and nondestructive updating processes is not easy to phrase in terms of interference theory. The concept of *unlearning* (Melton & Irwin, 1941) is certainly closely related to the notion of destructive updating, but there is at least one basic difference. Both ideas account for retroactive interference in the same general way, but the unlearning idea has always been viewed as a suppression-type process analogous to experimental extinction. Individual prior associations are elicited, overtly or covertly, and unlearned through some kind of active process. Even in the case of the *response-set suppression* hypothesis of Postman, Stark, and Fraser (1968), where it is not assumed that individual responses are elicited and unlearned but, rather, that the entire set of responses appropriate to prior learning is suppressed, it is still assumed that the suppression is an active process directed at the to-be-suppressed information. In either case, it is assumed that much of what is suppressed will spontaneously recover with time. The destructive-updating notion, at least as phrased here, does not involve an active effort to suppress the out-of-date information. Rather, the "destruction" of past information is assumed to be a consequence of the active encoding of current information. The process is more of a masking-type process, and there is no intrinsic reason for the masked information to recover spontaneously with time (at least not in terms of absolute strength).

The idea of structural updating does not have a natural representation in terms of interference theory. Structural updating presumes that the current input of information together with the series of preceding inputs is amenable to a longitudinal organization of some kind. Typically, interference theorists have concerned themselves with the acquisition of successive paired-associate lists in which an individual paired associate, from the subject's standpoint at least, lends itself to acquisition as an independent item. There is no natural basis on which to organize individual paired associates, either within a list or across lists. There are, of course, exceptions, such as some of the mediation paradigms. Also, in a recent experiment, Postman and Gray (1977) included a condition that can be viewed as a kind of structural updating.

Postman and Gray contrasted two different kinds of second-list learning in the standard A-B, A-D paired-associate transfer paradigm. In the "substitution" condition, subjects learned the second list in the normal way, presumably substituting A-D associations for A-B associations. In the "accretion" condition, however, when subjects were presented with a given stimulus (A) during List 2 acquisition, they were required to recall the appropriate B response from List 1 as well as the D response from List 2. In the accretion condition, then, subjects presumably added the D response to the A-B structure they had generated during List 1 learning, which would make the accretion condition roughly analogous to the story-construction strategy in the R. A. Bjork and McClure (1973) experiment reported above. Postman and Gray found that the accretion condition, when contrasted with the substitution condition, produced better list identification of responses, improved second-list retention, and at the same time did not slow down the initial acquisition of List 2.

The presence or absence of within-list structure may also be an important factor in transfer and interference between lists. Christen and Bjork (1976) have found, for example, that when subjects memorized five successive lists of words using the method of loci, where the loci were the exact same geographical locations on the same imagined path for each list, the PI one might expect on the basis of research with traditional paired-associate lists was completely absent. From a formal standpoint, it would seem that Christen and Bjork's experimental situation corresponds to the A-B, A-C, A-D, . . . paradigm, in which there typically is potent PI. Bower and Reitman (1972) also failed to find PI in a quite similar experiment. Christen, in a recent unpublished experiment, has even failed to find PI with repeated use of the same loci when each list consisted of the same words, but where the assignment of particular words to particular loci was changed from list to list (i.e., the A-B, $A\text{-}B_r$ paradigm). Whether the absence of PI in these experiments has something to do with the within-list structure provided by the connected loci, as I suspect, or whether the retention intervals employed were simply not long enough to permit the recovery of PI, is not clear. In any case, such experiments demonstrate a potential vulnerability of interference theory. At the same time that it is the most comprehensive theory in the field of human learning and memory, it has been based on phenomena from such a restricted set of materials and paradigms that its generalization to other materials and paradigms may be quite limited.

As far as the role of environmental context in updating, or the failure thereof, the interpretation offered above is very similar to *stimulus generalization decrement* and equivalent conceptualizations offered by interference theorists since the time of McGeoch (1942).

B. REGRESSION, RECOVERY, AND REINSTATEMENT

In terms of interference theory, the breakdown of updating corresponds to the recovery of PI (see, e.g., Postman, 1961; Underwood, 1957). Outdated (List 1) responses that are unlearned or suppressed during updating (List 2 learning) spontaneously recover with time. Postman and Underwood (1973) no longer feel that such a mechanism fully explains the rapid forgetting of List 2 responses, but the recovery postulate remains a central feature of interference theory.

As pointed out earlier, the regression process, which I presume to be a basic property of human memory, could take a form that is analogous to unlearning plus recovery. The possible alternative form of the regression process—regression to an abstracted prototype—is not an idea that has a straightforward representation in terms of interference theory.

The breakdown of updating owing to reinstatement (or reexposure) of to-be-forgotten items might be interpreted in interference theory as a loss of list differentiation. If to-be-forgotten items are retrieved or represented during or after the acquisition of current items, temporal and other bases for the differentiation of to-be-remembered and to-be-forgotten information are blurred. Loss of list differentiation has long been advocated by Underwood and his co-workers (e.g., Thune & Underwood, 1943; Underwood & Ekstrand, 1966) as an important factor in PI. Bennett's (1975) systematic analysis of PI in the Brown–Peterson paradigm is also congenial to a retrieval discriminability interpretation.

In my view, the role of reinstatement in PI has been underestimated and the role of spontaneous recovery has been overestimated. Under conditions that disincline subjects to retrieve or rehearse first-list responses during or subsequent to second-list learning, there may be no recovery of PI at all. Note that in the Bjork *et al.* data in Fig. 4, there was no recovery of PI in the FR condition [compared with the (−)R condition] when recall was delayed by arithmetic. In a similar condition in another experiment reported by Bjork *et al.*, there was also no recovery of PI. It is only when some of the to-be-forgotten items are presented again during the retention interval that a recovery of PI is observed.

In directed-forgetting paradigms (see R. A. Bjork, 1972, for a description of such paradigms) that are analogous to the traditional List 1, List 2 paradigms used by interference theorists, there is an explicit signal to subjects to forget the information presented prior to the signal. There is no reason for subjects to rehearse the to-be-forgotten items once the signal is presented because they do not expect to be responsible for those items. In the typical A-B, A-D paired-associate list learning experiment, on the other hand, there is no particular reason for subjects, from their standpoint, not to retrieve List 1 responses during List 2 learning or to rehearse List 1 responses during the retention interval following the acquisition of List 2. Such activities, as Houston (1966) has argued, could be the source of the observed recovery of PI. Houston has shown that if subjects do not expect a delayed test, there is no recovery of PI. Along the same lines, Postman and Gray (personal communication) have found that explicit instructions to subjects to give no thought to the first list during acquisition of the second list both speeded second-list acquisition and led to little or no recovery of PI. In general, at least until recently, interference theorists have been somewhat remiss in not concerning themselves more with the influence of control processes such as rehearsal and encoding strategy. That such processes can have a major impact on interference and transfer phenomena how seems undeniable. One need, for example, look no farther than the data in Fig. 2.

V. Concluding Comments

How do we keep, or fail to keep, our memories current is an important question from both a practical and a theoretical standpoint. In this chapter I have attempted to isolate some researchable aspects of that general question. As is clear from the preceding section, an analysis of updating not only raises some new and interesting issues, but also offers a new perspective on some old and fundamental issues. How far that new perspective will take us remains to be seen.

REFERENCES

Atkinson, R. C., & Shiffrin, R. M. Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2). New York: Academic Press, 1968. Pp. 89-105.

Bennett, R. W. Proactive interference in short-term memory: Fundamental forgetting processes. *Journal of Verbal Learning and Verbal Behavior,* 1975, **14,** 123-144.

Bjork, E. L., Bjork, R. A., & Glenberg, A. M. The reinstatement of interference owing to to-be-forgotten items. Paper presented at the meeting of the Psychonomic Society, St. Louis, November 1973.

Bjork, R. A. Theoretical implications of directed forgetting. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: Winston, 1972.

Bjork, R. A. Short-term storage: The ordered output of a central processor. In F. Restle, R. M. Shiffrin, N. J. Castellan, H. R. Lindeman, & D. B. Pisoni (Eds.), *Cognitive theory* (Vol. 1). Hillsdale, N.J.: Lawrence Erlbaum Associates, 1975.

Bjork, R. A., & McClure, P. Encoding to update one's memory. Invited paper, meeting of the Midwestern Psychological Association, Chicago, May 1974.

Bower, G. H., & Reitman, J. S. Mnemonic elaboration in multilist learning. *Journal of Verbal Learning and Verbal Behavior,* 1972, **11,** 478-485.

Christen, F., & Bjork, R. A. On updating the loci in the method of loci. Meeting of the Psychonomic Society, St. Louis, November 1976.

Craik, F. I. M., & Watkins, M. J. The role of rehearsal in short-term memory. *Journal of Verbal Learning and Verbal Behavior,* 1973, **12,** 599-607.

Houston, J. D. Proactive inhibition and undetected retention interval rehearsal. *Journal of Experimental Psychology,* 1969, **82,** 511-514.

Hunt, E., & Love, T. How good can memory be? In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: Winston, 1972. Pp. 237-260.

McGeoch, J. A. *The psychology of human learning.* New York: Longmans, Green, 1942.

Melton, A. W., & Irwin, J. M. The influence of degree of interpolated learning on retroactive inhibition and the overt transfer of specific responses. *American Journal of Psychology,* 1941, **54,** 157-173.

Postman, L. The present status of interference theory. In C. N. Cofer (Ed.), *Verbal learning and verbal behavior.* New York: McGraw-Hill, 1961. Pp. 152-179.

Postman, L., & Gray, W. Maintenance of prior associations and proactive inhibition. *Journal of Experimental Psychology: Human Learning and Memory,* 1977, **3,** 255-263.

Postman, L., Stark, K., & Fraser, J. Temporal changes in interference. *Journal of Verbal Learning and Verbal Behavior,* 1968, **7,** 672-694.

Postman, L., & Underwood, B. J. Critical issues in interference theory. *Memory & Cognition,* 1973, **1,** 19-40.

Smith, S., Glenberg, A. M., & Bjork, R. A. The influence of environmental context on recall and recognition. *Memory & Cognition,* 1978, in press.

Thune, L. E., & Underwood, B. J. Retroactive inhibition as a function of degree of interpolated learning. *Journal of Experimental Psychology,* 1943, **32,** 185-200.

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

Underwood, B. J. Interference and forgetting. *Psychological Review,* 1957, **64,** 49-60.

Underwood, B. J., & Ekstrand, B. R. An analysis of some shortcomings in the interference theory of forgetting. *Psychological Review,* 1966, **73,** 540-549.

Woodward, A. E., Bjork, R. A., & Jongeward, R. H. Recall and recognition as a function of primary rehearsal. *Journal of Verbal Learning and Verbal Behavior,* 1973, **12,** 608-617.

Yntema, D. B., & Mueser, G. E. Remembering the present states of a number of variables. *Journal of Experimental Psychology,* 1960, **60,** 18-22.