

SOME SPECULATIONS ON STORAGE AND RETRIEVAL

PROCESSES IN LONG-TERM MEMORY

by

R. C. Atkinson and R. M. Shiffrin

TECHNICAL REPORT NO. 127

February 2, 1968

PSYCHOLOGY SERIES

Reproduction in Whole or in Part is Permitted for
any Purpose of the United States Government

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

Some Speculations on Storage and Retrieval
Processes in Long-Term Memory

R. C. Atkinson and R. M. Shiffrin

Stanford University
Stanford, California 94305

ABSTRACT

A brief outline of the memory system is followed by somewhat speculative proposals for storage and retrieval processes, with particular care being given to distinguishing structural components from control processes set up and directed by the subject. The memory trace is conceived of as an ensemble of information, possibly stored in many places. For a given set of incoming information, the questions dealt with are whether to store, how to store, and where to store; the last question in particular deals with storage along various dimensions. Retrieval consists of a search along storage dimensions utilizing available cues to limit the search area and provide appropriate entry points. Both storage and retrieval are considered to take place in two steps, one consisting of a highly directed process under control of the subject and the other consisting of a pseudo-random component.

Paper presented at a conference on "Research on Human Decision Making" sponsored by the NASA-Ames Research Center at Moffett Field, California, 31 January - 2 February, 1968.

This paper will take a fairly speculative look at the structure of long-term memory, at the storage and retrieval processes by which information is placed in and recovered from long-term memory, at the joint operation of the short- and long-term stores, and at the control processes governing these various mechanisms. While the discussion will be primarily theoretical with no attempt made to document our assumptions by recourse to the experimental literature, some selected experiments will be brought in as examples. We will begin by outlining the overall conception of the memory system, a conception which emphasizes the importance of control processes. Long-term storage and retrieval will then be discussed in terms of the basic assumption that stored information is not destroyed or erased over time. This assumption may of course be relaxed, but we employ it to demonstrate that forgetting phenomena can be satisfactorily explained by postulating that decrements in performance occur as a result of a decreasingly effective search of long-term memory.

The primary distinction in the overall system is between structural features of memory and control processes (Atkinson and Shiffrin, 1967). Structural features are permanent and include the physical structure and built-in processes that may not be varied. Examples are the various memory stores. Control processes, on the other hand, are selected, constructed, and modified at the option of the subject. The use of a particular control process at some time will depend upon such factors as the nature of the task, the instructions, and the subject's own history. Examples are coding techniques, rehearsal mechanisms, and certain kinds of search processes.

The main structural components of the system are the three major memory stores: the sensory register, the short-term store, and the long-term store. Each of these stores may be further subdivided on the basis of the sensory modality of the stored information; such evidence as is available indicates that memory processes may differ somewhat depending on the sense modality involved (Posner, 1966). The sensory register accepts incoming information and holds it fairly accurately for a very brief period of time; a good example is the brief visual image investigated by Sperling (1960) and others, which decays in several hundred milliseconds. The short-term store (STS) is the subject's working memory in that the various control processes are based in it and directed from it. Information is selectively entered into STS from both the sensory register and the long-term store (LTS) and will decay from this store in about 30 seconds, except for control processes (such as rehearsal) which permit the subject to maintain the information in STS as long as desired. The long-term store is a permanent repository for information, information which is transferred from STS.

PROCESSES IN LONG-TERM MEMORY

The remainder of this paper will deal primarily with LTS, and also with STS in its capacity for handling LTS storage and retrieval. It would now be appropriate to outline our theory of long-term memory and define the most important terms that will be used. Long-term memory processes are first divided into storage and retrieval processes. These two processes are similar in many ways, one mirroring the other. Storage consists of three primary mechanisms: transfer, placement and image

production. The transfer mechanism is based in the short-term store and includes those control processes and mechanisms by which the subject decides what to store, when to store, and how to store information in LTS. The placement mechanism determines where the ensemble of information under consideration will be stored in LTS. It in turn will consist of directed and random components. Having decided finally where to store the ensemble of information, the image production process determines what parts of that ensemble will be permanently stored in that location of LTS. In general, not all the information desired is stored, and conversely, some unwanted information may be stored. The final ensemble of information permanently stored in LTS is called the image. This image is assumed to remain intact over time and during storage of other information. Retrieval, like storage, consists of three primary mechanisms: search, recovery, and response generation. Search is the process by which an image is located in memory, and like placement, consists of directed and random components. Recovery is the process by which some or all of the information in a stored image is recovered and made available to the short-term store, and response generation consists of the processes by which the subject translates recovered information into a specific response. We shall now turn to a detailed consideration of each of the processes outlined above.

Storage: Transfer

Transfer refers to the mechanisms by which information that has entered STS is manipulated there prior to placement in the long-term store. These mechanisms include a number of control-processes having to do with deciding what information to attempt to store, when to

attempt the storage operations, and what form of coding or other storage procedure should be employed. Before describing these control processes further, it should be pointed out that transfer involves at least one unvarying structural characteristic: whenever any information resides in the short-term store, some transfer of this information can take place to long-term store. The strongest evidence for this comes from studies of incidental learning (Saltzman and Atkinson, 1954), and from experiments first carried out by Hebb (1961) and Melton (1963). In these latter experiments subjects are given a series of digit spans to perform: for each span the subject is required to repeat back in order a short sequence of digits just presented. Unknown to the subject, a particular sequence is repeated at spaced intervals. Performance on the repeated sequence improves over trials, indicating that information about that sequence is being stored in LTS, even though the nature of the task is such that the subject does not attempt to store information about the individual spans in LTS. This assumption, of course, implies that images are being stored not only during "study" periods, but whenever information is input to the short-term store: during test, during rest periods, during day dreaming, and so forth. (Most laboratory experiments are designed to insure that essentially all storage takes place during study periods, but this is not always the case.)

In many situations, especially the typical experimental paradigms, a large amount of information is being input sequentially to the short-term store. In such a situation, the short-term store will act as a time-sharing system and the subject will select some subset of the presented information for special processing in STS such as rehearsal

or coding. The information not given special attention will decay and be lost from STS fairly quickly; LTS storage of this information will therefore be weak and undirected. If information is maintained in STS via simple rehearsal, but no special storage procedure such as coding is used, then the LTS image will be stronger than in the absence of rehearsal, but its placement will be quite undirected and thus the item will be difficult to retrieve at test (see Atkinson and Shiffrin, 1967). The selection of particular items for active attempts at storage will depend upon a number of factors. Items already felt to be retrievable from LTS will be dropped from active consideration; time would be better spent storing new, unknown information. There are many storage strategies the subject can adopt which result in the selection of particular items for processing: for example, in a paired-associate experiment with all responses being either X or Y, the subject might decide to store only the associates with the response X and to guess Y as a response to any unknown stimulus at test. Differential payoffs can also induce selection: items with higher payoffs being selected for storage. This phenomenon is illustrated in studies of reward magnitudes (Hurley, 1965). If two separate lists contain items with different payoffs, performance does not differ between the lists. If items within a list have different payoffs, however, the items worth more are preferentially selected and performance is better for them. Finally, in experiments where no great demand is made on the short-term system, all items can be given special storage procedures even if there is no need to do so.

What to transfer is dependent not only on the items presented for study, but also upon varying strategies the subject may adopt. Thus

the subject may attempt to cluster several items currently in STS and store them together. This obviously occurs in serial learning tasks, and often in free-verbal recall. Sometimes all the information in the presented item is not necessary for correct responding; in these cases the subject may decide to store only the relevant characteristics of the input. Most often the subject will select relevant characteristics of the input and then add to this information other information from LTS. In coding a paired-associate for example, the subject may recover a mediator from LTS and then attempt to store the paired-associate plus mediator. Note that the ensemble of information that the subject attempts to store and the ensemble that is actually placed in LTS are by no means identical; the latter may contain a large amount of information that the subject would regard as "incidental" or useless.

How to store the selected information refers largely to the control process adopted. In most cases a consistent strategy will be adopted and used throughout an experiment. These strategies include rehearsal, mnemonics, imagery, and other forms of coding. The level of performance will be greatly affected by the strategy used, the reasons for this becoming evident later in the paper.

Storage: Placement

Placement and search are two processes that have received little systematic consideration in the memory literature but are nevertheless extremely important. Placement refers to where in LTS storage of a particular information ensemble is attempted. By "where" we do not refer to a physical location in the cortex, but to a position in the

organization of memory along various informational dimensions.* These dimensions include sensory characteristics of the input (e.g., visual, auditory, or tactile storage), meaningful categorizations such as noun vs verb, or animal vs vegetable, and other characteristics such as the syntactic and temporal aspects of an item. These and other dimensions of storage will be elaborated further in the succeeding discussion.

There are two components to the placement mechanism; these will be called directed and random. Directed refers to that component of the placement mechanism which is specified by the control processes the subject is using, the information ensemble being stored, and the subject's past history of placement. Given these same conditions at a later time, the directed component will direct placement to the same LTS location. Furthermore, the search process during retrieval can follow the directed component to the same area of LTS. The second component of placement is random; it will occur as a result of local factors which change from one moment to the next and can be regarded as essentially random in nature. Thus at certain branches in the placement processes a succeeding storage attempt might select at random a different memory dimension and multiple stored images of the same information

* Anatomical evidence such as the Hubel and Wiesel (1962) explorations of information abstraction in the visual cortex of the cat, or the work of Penfield and Roberts (1959), or the older work on motor areas of the cortex, suggests that there may be a topographic placement mechanism. If one is trying to use a visual image to store a noun-noun pair (rather than, say, an auditory-verbal code) it would not be surprising if storage took place roughly in the area of the visual cortex. However, the form of the correspondence of the subject's informational organization of LTS with the physical structure of the nervous system is tangential to the discussion of this paper.

ensemble could result. Furthermore, during retrieval each of the random branches of placement would have to be explored via search in order to locate the stored image.

Note that the directed-random distinction is not the same as the structure-control process distinction; although random placement is not under the control of the subject, part of directed placement is also not under the subject's conscious direction. The directed component has three major determinants that will be considered in turn. The first is the kind of information in the item presented for study (and also in the ensemble selected for storage). Thus presentation in a free-recall task of a card with LION printed on it in black capital letters might lead to placement in locations determined by any or all of the dimensions: black, capitals, letters, words, animals, printed words, and so forth. In this free-recall example, as in other situations, certain storage locations will be more effective than others; storage in an "animal" location is not effective if at test the subject does not recall that he stored any words in the "animal" region. On the other hand, if the task was one of categorized free recall, in which there were a number of animals in the list to be recalled, then placement in an "animal" dimension might be very effective, especially since the first animal word recovered is likely to cause the subject to search in the "animal" region.

The second directed placement determinant is that induced by strategies the subject may select. If the strategy involves the formation of a natural language mediator for a paired associate, then the informational content and origin of the mediator may indicate placement

dimensions for storage of the pair plus mediator, perhaps in the "natural language" area. On the other hand, the formation of a visual image for coding purposes might lead to placement in the "visual area." If a cohesive strategy is used which encompasses many items, (for example, the placing of coded paired associates in the successive rooms of an imaginary house), then the placement of different items might be directed roughly to the same location.

The third placement determinant is that induced by the subject's pre-existing organizational structure and history of placement of similar information in the past. This kind of placement may often occur not under conscious control of the subject, but may nevertheless be consistent over trials. These three determinants of directed placement are necessary in order that the subject may be able to "retrace" his path and find a stored image during retrieval and search.

Either at the will of the subject or not, placement of an information ensemble may occur in more than one location in LTS. For example, the subject may encode an associate in two different ways and then store both resulting codes in each of the two locations defined by the codes. Multiple placement of this kind is said to result in multiple images or multiple copies in LTS. The extent to which multiple placement occurs in the usual experimental tasks is open to question. In some tasks, such as those in which the one-element model has been applied successfully (Bower, 1961), it would appear that a single copy assumption best fits the data. Even in these cases, however, the multiple copy models may be applied if the very first copy stored is always capable of allowing a correct response; in this case the effects of multiple storage

are not observable if only correct and incorrect response data are recorded.*

It is too much to ask of a memory system that placement be entirely directed. This would be akin to a library with a complete and accurate filing system, but there are a number of reasons why such a high accuracy system would be unfeasible for the type of memory system outlined here. These reasons include the drastic consequences of small failures in such a system, and considerations of access times. Furthermore, we are assuming that placement and search are parallel processes and there is evidence that search processes at times operate more or less randomly (see Atkinson and Shiffrin, 1965). Consequently we assume that there is a considerable component of placement which is also essentially random. That is, if placement were completely directed, there would be no reason for search to be random to any degree. (We shall consider random search processes later.) Sometimes part of the directed storage may be unavailable during retrieval; that portion of the placement is then essentially random since the subject must initiate a random search to find the right storage location.

Storage: Image Production

An ensemble of information having been placed at some location for storage, the image production process determines what portion of this

* A number of interhemispheric animal studies (Sperry, 1961) have indicated that at least two copies are normally made, one in each hemisphere, but this may not involve placement. Rather, it seems that once an image has been produced, the corpus callosum is involved in an after-the-fact transfer of the image to the other hemisphere.

information is permanently stored as an image there. We cannot say much about this process except that it occurs in some partial or probabilistic manner: at test, subjects can often recall incidental material which is correct but irrelevant, even when the required answer cannot be recalled. Actually it is difficult to separate the effects of image production from those of its retrieval counterpart, recovery. Recovery refers to the extraction of information from a stored image which has been located. A conceivable method for separating these processes is based on the fact that it is sometimes possible to use cueing to elicit from a stored image information not recoverable in a first attempt.

We next consider the contents of the image: the range and form of the stored information. A single image may contain a wide variety of information including characteristics of the item presented for study (its sound, meaning, color, size, shape, position, etc.) and characteristics added by the subject (such as codes, mnemonics, mediators, images, associations, etc.). In addition, an image most probably contains links to other images (other information which was in the short-term store at the same time); these links can be regarded as a set of directions to the locations of related images in LTS. There is some question as to whether temporal information in the form of some sort of internal clock reading may be part of the image. It is our feeling that the ability to make temporal discriminations can be explained on the basis of contextual information and counting processes, rather than on the basis of a clock reading recorded on the image.

We make the assumption that images are essentially permanent; they do not decay or disintegrate over time given an intact, physiologically

normal organism. This assumption is made for simplicity. We feel it is possible to propose appropriate search and storage mechanisms that explain decreases in performance over time. Some ways in which this may be done will be suggested when the outline of the system is completed.

Retrieval: Search

At test the subject is given certain cues specifying the nature and form of the required response. Assume that the information necessary to generate a response is not at that time in the short-term store. The subject will then attempt to locate the relevant image, or images, in LTS. This attempt is called the search process. The search will be monitored by the short-term store. That is, at any moment the short-term store will contain a limited amount of information such as the search strategy being employed, part of the information recovered so far in the search, what locations in LTS have been examined already, and some of the links to other images that have been noted in the search but not yet examined. The short-term store will thus act as a "window" upon LTS, allowing the subject to deal sequentially with a manageable amount of information. In addition to the directed search monitored by STS there is a random, diffuse component engendered by the information currently in STS. Thus when, say, the stimulus member of a paired-associate is presented for test, it will enter STS and at once a diffuse search is initiated by this member: as a result a number of images will be activated including many of the associates of this stimulus. There will be feedback such that activated images will be entered into STS, but this must be quite selective since STS has only a limited search capacity. Thus many activated images, possibly including the desired

image, may not gain access to STS. As the search continues and new information enters STS, the diffuse pseudo-random search component will be re-elicited by the new STS information. Hopefully, a relevant image will eventually enter STS and be recognized as such.

As the above discussion has tried to indicate, there are directed and random components to the search process. The subject has a considerable amount of control over the directed component and we now consider this in some detail. As was true in placement there are three primary determinants of directed search. Search may first be directed by cues and characteristics of the information presented for test. Thus if "kaq" is presented as a test on a previously studied paired-associate, "kaq-cen," then search might be initiated along dimensions of things sounding like kaq, of words beginning with k, of nonsensical three letter combinations, and so on. On a free-recall test, search might be directed to the "most recent list of items." Secondly, search may be directed by strategies adopted by the subject. Thus a search for natural-language-mediators may be initiated following the presentation of a stimulus member of a paired-associate for test. Or perhaps a search is initiated in the region of visual images containing this stimulus member. One search strategy often used employs ordering of the search. For example, we are likely to do better when asked to name all 50 states if we search memory in an ordered fashion, say alphabetically or geographically, rather than in a haphazard fashion. Thirdly, search may be directed by historical patterns of search behavior that the subject has developed through consistent use.

In any event, to the extent that the subject can remember, he will (or should) attempt to utilize the same directed search strategy as the directed placement used during storage. If the subject stored a paired-associate via a visual image, it would clearly not be effective to search for natural language mediators at test. This provides a strong reason for a subject to utilize a single, consistent storage strategy during training, even though switching coding techniques from item to item might minimize "interference" and confusion.

In carrying out a directed search, information will be recovered from various images and placed in STS. If this information appears to be promising, perhaps in terms of its similarity to the test information, then the search may be continued in the same area and direction, either in terms of the dimensions being searched, or in terms of the links recovered from successive images. Thus the search may be visualized as a branching process with random and directed jumps. At some point it may be decided that a wrong location has been reached (a wrong branch examined); at this time the subject may return to an earlier location or branch if its whereabouts is still held in the short-term monitor. If not, a return may be made to the original test stimulus in order to restart the search.

A decision that is very important in the retrieval process concerns when to terminate an unsuccessful search; after all, the desired information may never have been stored in LTS. A number of termination rules may be adopted. In cases where the response period is restricted, the search may be terminated by the time limit. In other cases, an internal time limit may be set which, if exceeded, terminates the search.

It is likely that this internal time limit will be dependent upon the kind of information actually recovered; if this information seems relevant or close then the search may be extended considerably. Another criterion for termination might be successive search attempts ending at the same unproductive location in LTS. In some cases termination for this reason is used as a positive approach: most of us have sometimes experienced the feeling that "if I only stop thinking about it for a while I'll remember it." In certain tasks other termination rules will sometimes be applicable. In free recall, for example, a series of words is read to the subject who then tries to recall them in any order. During retrieval the subject may find that successive searches result in recovery of words already recalled; in this case a termination rule might be based on the number of successive recoveries of words already recovered.

Of equal importance to the termination rule for an unsuccessful search is the termination rule for a "successful" search. That is, it will often happen that partial or incomplete information is recovered such that the subject is uncertain whether a particular response is appropriate. Similarly, some portion of the response might be recovered and a decision must be made whether to continue the search for the remainder, or to guess based on the partial information. Decisions in this case are probably based on available response time, payoffs for correct or fast responses, probability of correct guessing, and so forth. Termination criterion of this sort are closely related to the response production process which will be considered shortly.

Retrieval: Recovery

Once an image has been located, it is appropriate to ask what information contained in the image will be entered into the short-term store. This process is called recovery. To an extent, recovery of part or all of the stored information will be probabilistic, depending upon such factors as the current noise level in the system. Furthermore, as noted earlier, since the short-term monitor is limited and selective not all recoverable information will be entered into STS. This problem will tend to arise in fast large-scale random searches, in which large amounts of information may be activated with relatively little of this information being relevant. Thus in any particular situation the recovery of all the information in a stored image is by no means certain. The recovery process could conceivably be isolated from the others outlined so far by utilizing various cueing conditions at test to try and make more and more of the stored information available.

Retrieval: Response Generation

Having terminated the search and recovered information from LTS, the subject is faced with the task of translating this information into the desired response. Actually, a fair amount of experimental work has examined this aspect of retrieval and our remarks here will not be particularly novel. It should be pointed out first that when we speak of recovery of information we do not imply that this information will be verbalizable or directly available in the conscious experience of the subject. In some cases partial information may result in nothing more concrete than a feeling of familiarity on the part of the subject. Thus, in many cases this aspect of the subject's performance might be well

represented by a decision-theoretic model in which the subject is attempting to filter information through a noisy background (e.g., see Wickelgren and Norman, 1966; Bernbach, 1967; Kintsch, 1967). A good part of the response generation process consists of what can be called the guessing strategy. In general, guessing refers to the subject's selection of a response on the basis of partial information. There are a large number of guessing strategies that can be adopted and they will not be considered in detail here. It should be realized, however, that the probability of a correct response may not always be related in an obvious way to the amount of information recovered; guessing strategies can complicate matters. For example, in a paired-associate experiment where a list of stimuli is mapped on to two responses X and Y, the subject may store only information about stimuli with response X and then always guess response Y when a stimulus is tested for which no information can be retrieved. In this case, no information will be recovered about Y pairs, but they will always be responded to correctly. This serves to emphasize again the importance of control processes in even the simplest experiments.

DISCUSSION

We have now traced information from its presentation through storage, retrieval and output. We have not described ways in which performance will decline with time and intervening items. One way in which this can occur involves the storage of an increasing number of images, without a corresponding increase in the accuracy of the placement and search processes. In order to illustrate this point, and also

indicate how the system may be applied in an actual situation, we may consider free-verbal recall. A number of lists of words are read to a subject. Following each list the subject attempts to recall as many of the words in the preceding list as possible, in any order. Two results of interest here are the facts that there are almost no intrusions from preceding lists, and that performance decreases as list length increases (Murdock, 1962). These effects are found even if short-term storage is obliterated (Postman and Phillips, 1965; Atkinson and Shiffrin, 1965), so we shall consider this experiment only from the point of view of LTS. One interpretation of the lack of intrusions would hold that the placement process directs information about successive lists to separate locations in LTS, and at test a directed search is made only of the most recent location. Let us assume that within a list, information about individual words is stored in a non-directed fashion in that list location. Call the amount of information stored for the i^{th} word, S_i . Then the amount of information stored altogether in the most recent list location will be $\sum S_i = S$. At test the search process is immediately directed to the most recent list location, but the search is random within that area. Assume that n random searches are made in this area during the time allotted for responding. By random search we mean that the probability of finding an image relevant to word i on a search will be S_i/S . The probability of recovering information from that image and then generating the correct word will depend of course upon the amount of information, S_i . Suppose that performance is the result of n independent random searches of this kind. What then will happen to performance as list length increases? S_i will remain the same but

$\sum S_i = S$ will increase. Since the probability of "hitting" any image on a search is S_i/S , this probability will decrease with an increase in list length. Thus decreases in performance with increasing list length can be explained with reference to problems inherent in the storage and retrieval processes, without the necessity of assuming loss of information from stored images.

This free-recall model has been applied successfully to a large amount of data (Atkinson and Shiffrin, 1967). The model is particularly interesting because it utilizes all three retrieval processes outlined in this paper. The directed search refers to location of the most recent list. A random search is then made within that list location. Images identified in the search may or may not have information recovered from them. The amount of information recovered then determines the probability of correct response generation.

The free-recall model is one possible application of the system described in this paper. Despite its relative success, the assumption that placement is random within a list location is probably only roughly correct at best. Certainly most subjects tie together some of the words within a list (Mandler, 1967; Tulving, 1962). Furthermore, the search itself may not be nearly as random as was assumed. A situation in which these possibilities are accentuated is that of categorized free recall (Cohen, 1963). In this type of experiment a number of the words within a single list fall into well-known categories (e.g., months of the year, numbers from 0 - 9, kinds of monkeys, etc.). In this case we would probably expect both placement and search to be directed down to the level of the category, rather than the level of the list. A model

which seems to work well for this type of task assumes that the initial search is random within a list location, but once one member of a category is reported a directed search is made through the other members of the category, with any presented item in the category having a constant probability c of being recovered.

Another question we might consider in our framework is the source of differences in performance between recognition and recall procedures. One primary source arises in the response generation process: the recovery of partial information in the search will lead to better performance in recognition than in recall. For example, being able to recover the first letter of a response may guarantee perfect performance on a recognition test, but virtually chance responding for recall.

Another source found in paired-associate tasks is related to the search process: recall provides only one member of the pair, and location of the stored image must be based on cues provided by this single member. In recognition, however, both a stimulus and a response member are presented and search for the relevant image in LTS may be based on cues provided by either or both members. Finally, another source of difference between performance in recall and recognition may be found in the storage process: expectation of a recognition test may allow easier storage than expectation of a recall test. That is, less detailed information would need to be stored about an item if the tests were recognition rather than recall. This might permit storage of items that would otherwise have been ejected from STS for lack of time to deal with them. One test of storage versus retrieval effects was carried out by Freund, Brelsford, and Atkinson (1967). At study a

paired-associate item was presented and the subject was told he was either going to be tested by recall, by recognition, or he was not told which form of test would be used. Comparison of performance for the four types of items (told recall-tested recall, told recognition-tested recognition, not told-tested recall, or not told-tested recognition) allows storage and retrieval effects to be separated. Using this design it was established that differences between recognition and recall depended on differences in retrieval and not on storage. However, it seems clear that the results depended upon the specific stimulus materials used; with appropriate stimulus materials storage differences might also be detected.

It is sometimes implicitly assumed by memory theorists that recognition tests (yes-no or old-new tests in the simplest cases) eliminate retrieval effects and that differences between the various recognition procedures may therefore be attributed to storage. This assumption would be most parsimonious if true, but there is insufficient evidence to justify it. From our viewpoint there is reason to assume that retrieval effects are not eliminated by using recognition tests. In some recognition tasks it is clear that search effects are present. For example, if a paired associate is presented and the subject is asked whether the correct response is being displayed with the stimulus, one procedure the subject will use is to search memory, find the correct response, and compare it with the one presented. Thus, even in the simplest cases it is likely that recognition involves a variety of retrieval and search processes. In this regard we can point to several factors which might favor recall over recognition tests. The recognition

condition may cause a premature termination of the search process because the subject thinks he can correctly identify a given response, while an extended search would recover the correct one. In a recognition task where an incorrect response alternative is displayed, the incorrect alternative may initiate inappropriate search patterns that consume time and otherwise hinder performance.

The above discussions illustrate one of the benefits of introducing a highly structured, albeit speculative, long-term memory system. Such a system can be quite productive of alternative explanations for a wide range of memory phenomena that less structured systems may not deal with effectively. This in turn leads to experiments designed to determine which explanations are applicable in which situations. It is unfortunately beyond the scope of this paper to apply the system to the many experimental results in long-term memory. Nevertheless, we hope that it has been of some value to outline the theoretical system. Parts of the theory have been incorporated in models for a variety of experiments (Atkinson and Shiffrin, 1965, 1967) but the overall framework has not previously been elaborated.

In this paper no attempt was made to compare our system with extant theories of long-term memory. Most of the current theories have been presented at a somewhat more general level than was used here, and the present system may therefore be liberally interpreted as an extension and elaboration of certain ideas already in the literature.

References

- Atkinson, R. C., and Shiffrin, R. M. Mathematical models for memory and learning. Technical Report 79, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1965. (To be published in D. P. Kimble (Ed.), Proceedings of the third conference on learning, remembering and forgetting. New York: New York Academy of Science.)
- Atkinson, R. C., and Shiffrin, R. M. Human memory: A proposed system and its control processes. Technical Report 110, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1967. (To be published in K. W. Spence and J. T. Spence (Eds.), The psychology of learning and motivation: Advances in research and theory, Vol. 2. New York: Academic Press.)
- Bernbach, H. A. Decision processes in memory. Psychological Review, 1967, 74, 462-480.
- Bower, G. H. Applications of a model to paired-associate learning. Psychometrika, 1961, 26, 255-280.
- Cohen, B. H. Recall of categorized word lists. Journal of Experimental Psychology, 1963, 66, 227-234.
- Freund, R., Brelsford, J. W., Jr., and Atkinson, R. C. Recognition versus recall: Storage or retrieval differences? Paper presented at the Psychonomic Society Annual Meeting, Chicago, October, 1967. (Mimeographed copies available.)
- Harley, W. F., Jr. The effect of monetary incentive in paired-associate learning using an absolute method. Psychonomic Science, 1965, 3 141-142.
- Hebb, D. O. Distinctive features of learning in the higher animal. In J. F. Delafresnaye (Ed.), Brain mechanisms and learning. London and New York: Oxford University Press, 1961. Pp. 37-46.
- Hubel, D. H., and Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. Journal of Physiology, 1962, 160, 106-154.
- Kintsch, W. Memory and decision aspects of recognition learning. Psychological Review, 1967, 74, 496-504.
- Mandler, G. Organization and memory. In K. W. Spence and J. T. Spence (Eds.), The psychology of learning and motivation: Advances in research and theory, Vol. 1. New York: Academic Press, 1967. Pp. 328-372.

- Melton, A. W. Implications of short-term memory for a general theory of memory. Journal of Verbal Learning and Verbal Behavior, 1963, 2, 1-21.
- Murdoch, B. B., Jr. The serial position effect of free recall. Journal of Experimental Psychology, 1962, 64, 482-488.
- Penfield, W., and Roberts, L. Speech and brain mechanisms. Princeton, N. J.: Princeton University Press, 1959.
- Posner, M. I. Components of skilled performance. Science, 1966, 152, 1712-1718.
- Postman, L. and Phillips, L. W. Short-term temporal changes in free recall. Quarterly Journal of Experimental Psychology, 1965, 17, 132-138.
- Saltzman, I. J., and Atkinson, R. L. Comparison of incidental and intentional learning after different numbers of stimulus presentations, American Journal of Psychology, 1954, 67, 521-524.
- Sperling, G. The information available in brief visual presentations. Psychology Monograph, 1960, 74 (Whole No. 498).
- Sperry, R. W. Cerebral organization and behavior. Science, 1961, 133, 1749-1757.
- Tulving, E. Subjective organization in free recall of "unrelated" words. Psychology Review, 1962, 69, 344-354.
- Wickelgren, W. A., and Norman, D. A. Strength models and serial position in short-term recognition memory. Journal of Mathematical Psychology, 1966, 3, 316-347.