Endel Tulving

Cue-Dependent Forgetting

When we forget something we once knew, it does not necessarily mean that the memory trace has been lost; it may only be inaccessible

We trail behind us, unawares, the whole of our past; but our memory pours into the present only the odd recollection or two that in some way complete our present situation.—Henri Bergson, *Creative Evolution*

Memory for an event is always a product of information from two sources. The first is conceptualized as the memory trace—information laid down and retained in a person's memory store as a result of the original perception of the event. Its postulation is necessary to account for the residual effects of the event. The other source is the retrieval cue—information present in the individual's cognitive environment at the time retrieval occurs. This construct is necessary to account for the high degree of selectivity exhibited by the retrieval system; from among a myriad of traces only a small subset is actively involved in producing any particular memory.

Endel Tulving, who was born in Estonia, received his B.A. from the University of Toronto and Ph.D. from Harvard. After 1956 he taught at the University of Toronto, moved to Yale in 1970, and is now Professor of Psychology at both of these institutions. His research has been focused on the interrelation between storage and retrieval processes in human memory. This article, which was written during his residence as a Fellow at the Center for Advanced Study in the Behavioral Sciences at Stanford in 1972–73, is an updated version of a paper read at the 19th International Congress of Psychology in London in 1969. Dr. Tulving's research described herein has been supported by the National Science Foundation and the National Research Council of Canada, and he is grateful to Miriam Gallagher and Gordon H. Bower for helpful comments. Address: Department of Psychology, Yale University, New Haven, CT 06520.

In its dependence upon simultaneous fulfillment of the two necessary conditions, remembering is like many other familiar psychological phenomena. We see a star in the sky if its light reaches our eyes *and* if there is no light reflected from the sky around it. We understand a spoken message if we hear a certain set of sound stimuli *and* if we know the language in which it is spoken. Similarly we remember an event if it has left behind a trace *and* if something reminds us of it.

To understand memory requires understanding the informational properties of memory traces and retrieval cues, the interaction between the trace system and the retrieval system, and the processes by which information from the two sources is combined and converted into a memory experience. In the quest of this understanding many questions are raised and many phenomena analyzed in laboratory experiments. An enduring problem has to do with the fundamental conditions responsible for forgetting. Forgetting, as defined here, is the inability to recall something now that could be recalled on an earlier occasion.

Experimental analysis of forgetting in the psychological laboratory is almost a hundred years old, but philosophical speculation about its nature covers a much longer time span. Throughout this long history, two general hypotheses have provided two different vantage points for conceptual analyses of forgetting. One holds that forgetting is a trace-dependent phenomenon which occurs because certain changes take place in the specific traces of events. The traces decay, or deteriorate, or are somehow lost from the store, and the information they originally may have contained becomes unavailable.

The other idea is that forgetting is a cue-dependent phenomenon, reflecting the failure of retrieval of perfectly intact trace information. A person's cognitive environment will change over time, the relevant retrieval information is absent, and as a consequence the trace information *available* in the store becomes *inaccessible.* In addition to the advocates of these two rather different points of view, there are theorists who believe that forgetting may occur for either reason, or for both reasons together, as well as those who claim that retrieval failure is not "true" forgetting, and that "true" forgetting is only the failure of memory that comes about because of deleterious changes in the stored information.

In this paper I would like to present some evidence in favor of the hypothesis of forgetting as a cue-dependent phenomenon and argue that, at least within a limited domain, it is the preferred point of view. The evidence comes from rather simple experiments in which subjects—normal, intelligent adults—are asked to remember discrete events. A particular word (item) in a particular collection (list) of other words in a particular experimental situation (general context) constitutes a typical event about which information is stored in memory and subsequently retrieved. While each word in the list is known to the subject, the *occurrence* of the word in an experimental list constitutes a unique and original experience for him, and no trace of such an *event* exists in his memory store.

The use of verbal materials in these list-item experiments is simply a matter of convenience. Each word serves as a focal element of the to-be-remembered event. Its properties can be easily manipulated, the subject can readily discriminate it and label it in a predictable manner, and the subject's memory for each event can be reasonably safely inferred from his ability to produce or identify correctly the word in a subsequent test.

Constant traces—variable retrieval

The role that retrieval cues play in the act of remembering can be easily overlooked both by the rememberer and by someone who observes him, for instance an experimenter or a theorist of memory. People often seem to remember things "spontaneously," in absence of any particular instigators, and it is seductively simple to attribute such memories to the "strength" of their traces. Another reason for overlooking retrieval cues may lie in the fact that large variations can occur in recall of events under conditions where the retrieval situation is held constant, for instance when a "vivid" event is more readily recalled than a "dull" one, in the absence of any specific cues. This fact reinforces the idea that variations in recall reflect properties of memory traces quite independently of any properties of the retrieval environment. Thus we find even on the contemporary scene theories of memory that have nothing to say about retrieval cues or about the interaction between trace information and retrieval information—that *sine qua non* of all memory phenomena.

On the other hand, it is a relatively simple matter to demonstrate large effects on retrievability of stored information under conditions where the amount and organization of that information is held constant experimentally, and where only retrieval information is varied. In a recent experiment, Michael Watkins and I presented subjects with pairs of closely related words. Some pairs in the list consisted of strong associates (e.g. *bark-dog*), others of rhyming words (e.g. *worse-nurse*). Subjects had been told prior to the presentation of the list that the target items in this input list were the right-hand members of pairs and that their memory for these words would be tested. After seeing a list, subjects were given two successive tests. In the first, the left-hand members of pairs were provided as retrieval cues. Thus, for instance, the subject would be given the cue *bark* (associated with) ________, or *worse* (rhymes with) ________, and told to recall the corresponding word from the study list.

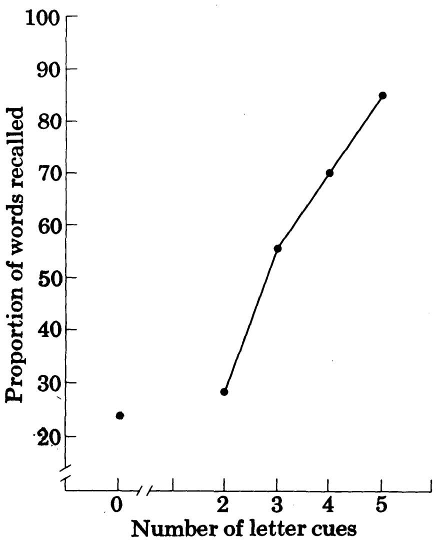

Figure 1. Proportion of five-letter words recalled from a list as a function of the number of initial letters of words given as cues. (Data from an experiment by Tulving and Watkins, in press.)

Subjects did quite well in this task, recalling 74% of associative target words and 56% of rhyming target words, but these data are of no particular interest. The more interesting data came from the second recall test, especially in respect to target words that subjects had failed to recall in the first test. Here the target word was cued with either an associative or rhyming word that had *not* appeared anywhere in the list, and the cue type was always different from the one used in the first test. Thus subjects would be given cues such as *grog* (rhymes with) ________, or *doctor* (associated with) ________, and they were expected to produce *dog* and *nurse*, regardless of whether they had recalled these words when the specific list cue was presented in the first test. The data relevant to this discussion was provided by probabilities of recall of target words cued by "new" cue words when the subject had *failed* to recall the same target words in the first test. The rhyming cues proved effective for 22%, and the associative cues for 30%, of the "forgotten" words. Although these probabilities are not high, they are considerably higher than zero.

It should be mentioned parenthetically that in this experiment, as well as in others we will discuss, observed effectiveness of cues cannot be explained in terms of guessing on the part of the subjects. Subjects usually refrain from guessing, and the amount of guessing that does occur is not sufficiently high to account for the observed effects of specific retrieval cues.

While there are other lessons to be learned from these data, the main point here is that, even when retrieval of a target event in the presence of the most "obvious" cue fails, there may be other cues that will provide access to the stored information. Failure of retrieval does not, therefore, necessarily imply weakening of associations between experimentally designated stimuli and responses, or loss of information from the traces of experienced events.

Consider next data from another simple experiment (Tulving and Watkins, in press). Subjects were shown lists of 28 five-letter words, and their recall was tested either with two, three, four, or five initial letters of these words as specific retrieval cues or without any specific cues at all. (In the case of five-letter cues, the cues were, of course, nominally identical with target words.) For instance, for the target word *grape*, different subjects in different groups would be given cues *gr---*, *gra--*, *grap-*, or *grape* and asked to write, beside the cue, the corresponding target word from the list they had studied. In tests where they were given no cues at all, they were simply asked to write, in any order, all the words they remembered from the list.

The principal data from the experiment are depicted in Figure 1, which shows a regular increase in the probability of recall as a function of the size of the word fragment given as cue. Since the to-be-remembered items were identical in all experimental conditions, as were the encoding operations performed by the subjects on the presented words, we know that the traces of target words in the store were identical for the different cueing conditions. Hence the data in Figure 1 illustrate how retrieval of information contained in a fixed memory trace depends on the amount of appropriate information present in the subject's cognitive environment at the time of the attempted retrieval. Since the level of retrieval varies systematically with the amount of retrieval information provided to him, it seems reasonable to conclude that even under the free-recall conditions, where no specific cues were provided, some retrieval information must have been available that made possible access to specific stored traces.

A third experiment (Light 1972) has demonstrated the effectiveness of retrieval cues containing information of yet another sort. On a single-input trial, subjects studied target words presented for memorization. The targets were presented either as isolated words, shown one at a time, or as parts of meaningful sentences, with one sentence shown at a time. Subjects had to recall the target words either without any cues or with one of three types of specific retrieval cues: (1) homonyms of target words, (2) synonyms of target words, or (3) cue words identical with target words. Since the subjects did not know how they would be tested, information stored was the same for all retrieval-cue conditions within a given presentation (encoding) condition, although it probably did vary between the two presentation conditions.

The results of this experiment are summarized in Table 1, which presents probability of recall in each of the four retrieval conditions for both encoding conditions. We see again that different retrieval cues showed large differences in their valence, or effectiveness. While in this situation the relative valence of different cues could not be specified in advance, the data can be interpreted in the same way as those shown in Figure 1: different kinds of retrieval cues are differentially effective in providing access to memory traces, and hence they can be assumed to vary in the amount of relevant retrieval information they contain.

Table 1. Proportion of words recalled from a list as a function of four types of retrieval cues combined with two encoding conditions (data from an experiment by Light 1972).

	Retrieval cues			
Encoding condition	*Cue words identical with target words*	*Cue words homonyms of target words*	*Cue words synonyms of target words*	*No cues*
Single word	.92	.81	.51	.32
Sentences	.92	.88	.40	.18

The findings from all three of these studies illustrate that recall of an event may fail solely because relevant information is absent from the retrieval environment. Since in these studies trace information was held constant experimentally, we can conclude that retrieval failure need not imply any impairment or deterioration of information contained in memory traces. The argument can be extended without any difficulty to the kind of retrieval failure that is referred to as forgetting—that is, failure of retrieval of an event in a test following successful retrieval of the same event in an earlier test. It is perfectly possible that, in this type of situation, too, impaired recall results from a lack of appropriate information in the retrieval environment rather than from a change in the specific stored information.

One implication of the view of forgetting as a cue-dependent phenomenon can be directly tested by manipulating the information in the retrieval environment after apparent forgetting of an event has taken place. If forgetting is a manifestation of loss of trace information, changes in the retrieval environment should be inconsequential. If, on the other hand, it reflects the absence of relevant retrieval information, appropriate changes in retrieval cues should produce reversal of "forgetting." To illustrate the effect of changes in the retrieval environment on recall of "forgotten" events, we will consider the results of three types of experiments: (1) a single-trial free recall task, (2) a multi-list retroaction paradigm, and (3) a subject-generated recognition task.

Forgetting in free recall

In a single-trial free-recall task the subject is shown a list of *n* familiar words once, and after its presentation he is asked to recall as many of these words as he can, in any order. In the first test, *T*1, after the presentation of each *individual word*, the subject is asked to "recall" what the word was. This test is usually not given because its results are highly predictable—subjects can always "recall" the word seen last! Test *T*2 is given at the end of the presentation of the list, and, typically, subjects fail to recall a certain proportion of the words. The question is, How can we account for such intra-trial forgetting?

A popular theory of memory, according to which the major components of the memory system are a short-term and a long-term store (e.g. Atkinson and Shiffrin 1968; Glanzer 1972; Waugh and Norman 1965), holds that every presented list word enters the short-term store, which is of rather limited storage capacity. As long as the word—more precisely, information about its occurrence in the particular list—resides in the short-term store, it (1) can be retrieved without any difficulty, (2) can be transferred into a more commodious long-term store, and (3) can be displaced by other incoming information. Thus, perfect recall at *T*1 simply reflects perfect accessibility of information in the short-term store, while forgetting observed at *T*2 is a consequence of the premature displacement of information from the short-term store, before its transfer into the long-term store has been completed.

The dual-store model of memory represents a theory of trace-dependent forgetting when applied to a single-trial free-recall task: the observed recall failure of words at *T*2 reflects the loss of information from the limited-capacity, short-term store. Thus, for instance, the relatively low level of retrieval under the free-recall (no cues) conditions in the Tulving and Watkins (in press) experiment depicted in Figure 1, and in the Light (1972) experiment presented in Table 1, would be interpreted by dual-store memory models in terms of the failure of the information to reach the long-term store.

In test *T*3 specific retrieval cues are presented for recalling the list material. The data from *T*3, showing recall to have increased over *T*2, are not readily compatible with the basic assumptions of dual-store models. For one thing, the assumption that intra-trial forgetting in free recall reflects failure of transfer of all list words into the long-term store makes it necessary to postulate a second mechanism of some kind to account for the enhanced recallability of the same information under changed conditions of retrieval. The hypothesis of cue-dependent forgetting, on the other hand, can be used to interpret the data from all three tests of retrieval in terms of a single major assumption—that the memory trace of each presented word is created extremely rapidly at the time the word is perceived and encoded and that access to it depends on retrieval information present.

Perfect recall at (imaginary) test *T*1 is attributable to the presence in the retrieval environment of information that is highly compatible with, or complementary to, that contained in the memory trace. At *T*2, the trace of each word still exists intact in the store, but the interpolated presentation of other words has changed the information in the retrieval environment so as to render recovery of information from some of the stored traces impossible. Presentation of specific retrieval cues at *T*3 restores to a large extent the retrieval information that was absent at *T*2, thus making possible the completion of the retrieval process.

In both Tulving and Watkins's and Light's experiments, tests *T*2 (free recall of the list, with no specific cues) and *T*3 (cued recall) were not separated in time. But this does not invalidate the argument for cue-dependent forgetting, particularly since it is simple enough to demonstrate that presentation of specific retrieval cues *following* noncued recall still has large facilitative effects on subjects' performance (e.g. Tulving and Pearlstone 1966). Thus it seems reasonable to attribute intra-trial forgetting in a single-trial free-recall task to the absence of appropriate retrieval information rather than to the absence of relevant trace information.

There are obviously several objections to this interpretation of intra-trial forgetting, but it is better to postpone their consideration until after we have looked at the data from the other two experimental situations.

Retroactive interference

An example of the second type of experiment is described by Tulving and Psotka (1971). Lists of 24 words were used: each list consisted of 4 words in each of 6 conceptual categories. For instance, one of the lists contained the following words: *hut, cottage, tent, hotel; cliff, river, hill, volcano; captain, corporal, sergeant, colonel; ant, wasp, beetle, mosquito; zinc, copper, aluminum, bronze; drill, saw, chisel, nail.* The words in the same category were always grouped together, as in the above example, to make quite clear to the subject what the conceptual categories were and to encourage him to encode the words accordingly. Different groups of subjects learned different numbers of such lists. Group 1 learned one list, Group 2 learned two, and so on, to Group 6, which learned six lists. All lists had the same overall structure, but both the categories and the words within them were different in the successive lists learned by a given group.

Each list was shown three times in succession, each time at the rate of one second per word; then a noncued recall test (*T*1) was given in which the subject had to produce as many words from the list as he remembered. After the subjects in a given group had learned their last list and had had the *T*1 test on it, they were asked to recall *all* the words from *all* the lists they had seen. In this test (*T*2), too, no specific hints or cues were given to the subjects. After this overall noncued test of all lists, subjects spent 10 minutes on an interpolated mental task that had no formal relation to the material they had studied, and then they were asked to recall the words once more, still in the absence of specific cues. Immediately after finishing this second overall noncued recall test (which we do not need to consider in the present analysis), the subjects were given a third overall test (*T*3). This time, subjects were presented with the names of all the conceptual categories they had derived from the lists studied. For instance, category labels for the sample list above—*types of building, earth formations, military titles, insects, metals, and carpenter's tools*—were listed on recall sheets, and the subjects were asked to write down all the words belonging to them that they had seen in their lists.

Figure 2 shows the mean number of words recalled from the 24-word lists in the three tests. The curve depicting the data from the second, noncued test, *T*2, describes the classical finding of retroactive interference in this type of experiment: the number of words "forgotten" from a list is directly related to the number of other lists interpolated between the learning of the list and *T*2 test.

The interference theory provides the most widely accepted interpretation of forgetting in this type of multi-list retroaction experiment. One current variant of the theory—the unlearning–recovery version—holds that recall losses in *T*2 reflect the unlearning of both specific (item-to-item) and general (context-to-item) associations in the list (Keppel 1968; Postman and Keppel 1967; Postman and Underwood 1973). Although it has never been clear to what extent the notion of unlearning of associations refers to the deterioration or degradation of the associations as stored information, the term *unlearning* does imply such degradation. On the other hand, unlearned associations are known to recover "spontaneously" in strength in the absence of any further interpolated learning, and the fact that such recovery occurs seems to be incompatible with the idea of permanent weakening of associations.

Another version of the interference theory—the response-set suppression theory—explains the greater forgetting of earlier rather than later lists in *T*2 as a consequence of generalized competition among response sets from different lists, or as a manifestation of the inertia of the selector mechanism (Postman and Stark 1969). While this theory implies that traces of items learned in earlier lists have not been lost from the store, unlike the cue-dependent hypothesis it assumes that recall is impossible because of competing responses rather than because of an absence of appropriate retrieval information.

Let us now return to the data from the third test in the experiment. The curve labeled "Cued recall" in Figure 2 depicts the subjects' level of recall of target words from all the lists when presented with the category names as specific retrieval cues. It shows that to a large extent these cues restored recall to its earlier levels. Thus in this type of multi-list retroaction experiment, observed forgetting is reversible.

The explanation of this reversal of forgetting requires no special assumptions if one accepts the hypothesis of cue-dependent forgetting. It does not, however, follow directly from either version of the interference theory. It is not clear how the presentation of specific retrieval cues could restore recall if impairment of recall is a consequence of unlearned associations. Nor is it immediately obvious why the presentation of retrieval cues should reduce general response competition or overcome the inertia of the selector mechanism, unless both of these mechanisms depend on the information available in the retrieval environment. And if this is so, the response-set suppression theory would become a special case of the general hypothesis of forgetting as a cue-dependent phenomenon.

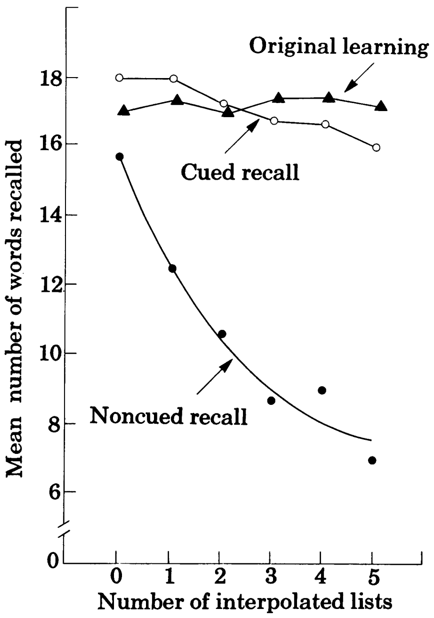

Figure 2. Mean number of words recalled from a list of 24 words in three successive tests—(*T*1) original learning, (*T*2) noncued recall, and (*T*3) cued recall—as a function of the number of other lists learned between the list and the second test (noncued recall). (Data from an experiment by Tulving and Psotka 1971.)

The interpretation of the data from the Tulving and Psotka experiment on the basis of the hypothesis of cue-dependent forgetting is very similar to that applied to the data demonstrating intra-trial forgetting in single-trial free recall. In the original learning test (*T*1), certain information is present in the retrieval environment that combines with the information in the specific memory traces to produce overt recall of target words. Retroactive interference observed in the overall noncued recall test (*T*2) reflects changes in the retrieval information, rather than loss of information from the memory traces. The changes in the informational content of the retrieval environment are brought about by interpolated learning and recall of other lists. The presentation of category names restores the missing information to the learner's cognitive environment at the time of test *T*3, thereby making possible recovery of the information laid down in the memory store at the time of initial learning.

Failure of recognition

The third set of data of interest concerns recognition memory. In a typical recognition-memory experiment the subject is shown a list of familiar words and asked to remember them. Then a test is given in which list words are mixed with new words, and the subject is asked to decide which words were on the previously seen list.

Many theorists believe that the problem of retrieval of stored information in a recognition memory test is simpler than that in a recall situation (for a review, see McCormack 1972). When faced with a test item, all the subject needs to do, so the argument goes, is to decide whether or not he saw it on the list. In the recall test, on the other hand, the subject must first implicitly generate a number of candidate words and then make a recognition decision about each (Bahrick 1970; Kintsch 1970). This dual-process theory of recall is a theory of trace-dependent forgetting when applied to the recognition situation, because it assumes that failure of recognition of an item learned earlier reflects absence of appropriate information in the memory store. There is no problem of retrieval in the recognition test—only a problem of decision. According to such a formulation, it makes little sense to talk about recognition memory in terms of an interaction between stored information and information present in the retrieval environment, because the latter is assumed to be held constant by the design of the experiment.

Since the dual-process theory is widely accepted, it is of particular interest to consider how the cue-dependent theory would handle forgetting in a recognition test. The illustrative data in this case come from the first of three experiments described by Tulving and Thomson (1973). A schematic description of the experiment is shown in Figure 3. All subjects received the same treatment sequence consisting of four steps. First, they were shown 24 pairs of words, such as *pretty*–BLUE, *glue*-CHAIR, and *grasp*–BABY. They were explicitly told that their task was to remember the second member of each pair.

Second, after the presentation of the list the subjects were provided with other words closely related to the to-be-remembered list words (e.g. *sky, table,* and *infant*) and

Step 1 Study list		Steps 2 and 3 Free-association and recognition test					Step 4 Cued-recall test	
head	LIGHT	dark	night	light	(black)	room	head	light
bath	NEED	want	need	desire	wish	get	bath	____
pretty	BLUE	sky	sun	cloud	blue	open	pretty	blue
grasp	BABY	infant	child	mother	love	(baby)	grasp	____
whiskey	WATER	lake	water	smooth	river	ocean	whiskey	water
cabbage	ROUND	square	(round)	flat	circle	corner	cabbage	round
spider	BIRD	eagle	eye	mountain	bird	high	spider	____
glue	CHAIR	table	chair	desk	lamp	top	glue	chair

Figure 3. Schematic outline of the experiment demonstrating failure of recognition of recallable words. In Step 1, subjects studied cue-target pairs with a view to remembering the target words, in capital letters. In Step 2, they wrote in words in response to strong extra-list associates of target words. In Step 3, subjects circled words they thought were target words from the study list. The hypothetical subject whose work is shown here produced 8 target words and circled 2 of them correctly (BABY and ROUND) as well as one incorrectly (BLACK), for a recognition score of 2 out of 8, or 25%. In Step 4, the list cues were presented to the subjects, and they were asked to recall the corresponding target words. The subject above correctly recalled 5 out of 8 words, or 62.5%, a typical result.

asked to write down several free associations to each. A typical subject generated four words in response to each stimulus word. The generated words included over 70% of the target words the subject had seen in the list. Thus, subjects frequently produced the word *blue* in response to the stimulus *sky, chair* to *table,* and *baby* to *infant*—these generated copies of target words were, of course, mixed in with words that had not occurred in the input list.

Third, subjects were told to look carefully at each word they had generated in the free-association test and to circle every word they recognized as a target word from the list. They were able to recognize correctly 24% of all the generated copies of the target word. Finally, as the fourth step in the experiment, subjects were given the cue words that had been paired with the target words at input—such as *pretty, glue,* and *grasp*—and asked to recall the corresponding target words from the list. In the presence of these retrieval cues, subjects could recall 63% of the target words. Thus, the level of retrieval was considerably higher in the presence of input cues which had appeared with the target words in the study list than in the presence of literal copies of target words in the subject-generated recognition test. Subjects could recall many list words that they did not recognize.

Since it is a safe assumption that any subject would have been able to recognize each target word immediately after its presentation in the list ("test" *T*1), the failure to recognize copies of many list words in the self-produced recognition test (*T*2) represents an instance of forgetting. From the point of view of most current theories of recognition memory, such forgetting would have to be regarded as reflecting some kind of deterioration of stored information. In particular, any theory that conceptualizes the recognition task as involving no retrieval problem and sees the subject as only having to make a decision about context information stored with an item would have to interpret recognition failure as direct evidence of loss of information about the list item in the memory store. According to these theories, recall following recognition failure would be impossible, since recall requires information whose existence is denied by recognition failure.

According to the hypothesis of cue-dependent forgetting, on the other hand, the recognition failure in test *T*2, like all other failures of retrieval of information about events and episodes, reflects absence of appropriate retrieval information rather than loss of information from the store. The subject-generated copies of target words apparently did not contain sufficient retrieval information, while list cues that had been paired with target words at input apparently did—which is why cued recall was higher than recognition. The precise reasons for such an outcome are not yet entirely clear, although several possibilities for explaining the data exist (Tulving and Thomson 1973); but this—we hope, temporary—gap in our knowledge does not invalidate the general reasoning underlying the hypothesis of cue-dependent forgetting. Thus, even failure of retrieval in a recognition situation need not mean that relevant memory traces are not available in the store; it may mean only that appropriate information is missing in the retrieval environment.

Advantages of the hypothesis

The hypothesis of cue-dependent forgetting has several advantages over those that attribute forgetting to deterioration of traces. Let us consider some of them briefly.

The first advantage is a matter of parsimony: one general form of the cue-dependent view can account for diverse instances of forgetting for which different versions of trace-dependent forgetting have been postulated. We have just seen how three different trace-dependent mechanisms were advanced to account for retrieval failure in three different memory tasks. Intra-trial forgetting in free recall was explained in terms of the displacement of information from the short-term store and consequent failure of its transfer into the long-term store. Retroactive interference in the multi-list free-recall task was attributed to unlearning of general or specific associations, or to suppression of learned response repertoires induced by generalized competition. And subjects' inability to recognize learned words in the self-produced recognition test was seen as evidence for the loss of list-specific information about the words from the store. A single hypothesis of cue-dependent forgetting, however, fitted the data in all three instances.

Second, the claim that observed forgetting in a particular situation is a consequence of loss of appropriate information from the retrieval environment can be—and, in the three cases we have considered, was—directly verified by restoring the lost information to the learner's cognitive environment. The presentation of category names of "forgotten" list words in both free recall and retroaction tasks, and the presentation of list cues in the experiment demonstrating recognition failure of recallable words, led to the recovery of many of these "forgotten" target words. The major claim of theories of trace-dependent forgetting—that information is somehow lost from the stored traces—is not amenable to the same kind of direct experimental verification.

The third advantage has to do with the fact that sometimes recall *increases* over a retention interval, in the absence of interpolated learning and under conditions where the two tests are nominally identical. The hypothesis of cue-dependent forgetting requires no additional assumptions to accommodate this phenomenon of reminiscence, as it is called in the technical literature, or the closely related phenomenon of spontaneous recovery that we have already mentioned. It is entirely conceivable that changes in the retrieval environment, which we know is at least partly determined by the person's informational intake and mental activity, can become more appropriate for retrieval as well as less appropriate. It is more difficult to imagine why traces should sometimes grow stronger and sometimes weaker over time. Trace-dependent theories would require additional assumptions to explain reminiscence and spontaneous recovery.

Fourth, our hypothesis of cue-dependent forgetting not only makes it meaningful to search for situations in which recall is "higher" than recognition, but also makes it possible to understand how this otherwise somewhat implausible phenomenon can come about. If remembering an event is a joint product of information contained in the trace and that contained in retrieval cues, then it is at least logically possible that, with trace information held constant, retrieval information contained in some other cue might better complement the trace information than the information contained in a particular literal copy of the original event. We now know that this logical possibility has a corresponding psychological reality under certain experimental conditions.

Finally, the hypothesis of cue-dependent forgetting leads to new questions about successful and unsuccessful remembering that are not so readily suggested by trace-dependent theories. The most important concern the act of retrieval as such. Precisely what is the nature of information contained in the retrieval environment that is necessary and sufficient for access to information contained in the memory trace? Why are some retrieval cues effective while others that, on intuitive or other pre-experimental grounds, seem to be equally potent turn out to be quite ineffective (e.g. Freund and Underwood 1969; Slamecka 1968)? How are we to conceptualize the nature of the interaction between stored information and information contained in retrieval cues? Research on these and similar questions will tell us something about both the processes and mechanisms of remembering and the reasons for occasional memory failure. It also holds out the promise that eventually an explanation of forgetting will be included in the theory of remembering and will not require special treatment.

Objections and criticisms

The hypothesis of cue-dependent forgetting represents only a general starting point of the analysis of memory phenomena subsumable under the heading of "forgetting," and there are many more specific problems that have to be worked out. There are also certain objections to and criticisms of the hypothesis that recur in the literature, and these deserve some comment.

Some people find it difficult, on a variety of grounds, to accept the assumption that stored information does not change over time. Strictly speaking, however, this is not an assumption that is entailed in any way in the hypothesis of cue-dependent forgetting. The hypothesis only claims that it is *not necessary* for a theorist to assume changes in memory traces to account for forgetting. It is both conceivable and plausible that traces do change over time and that information contained in them is lost, and it follows logically that such loss would be a sufficient condition of forgetting. But it is important to remember that no evidence exists to compel the conclusion that loss of trace information is responsible for any observed instance of forgetting. In some instances we can demonstrate that forgetting is reversible through the manipulation of the retrieval environment and that therefore in these cases forgetting would not have been a consequence of trace deterioration. In other cases, where reversibility has not been or cannot be demonstrated, it is no more reasonable to suppose that memory traces have changed than to suppose that the retrieval environment has changed.

Another criticism, closely related to the first one, has to do with the alleged assumption that the hypothesis of cue-dependent forgetting denies permanence of forgetting. Some critics have reasoned that, if all forgetting were a consequence only of the lack of appropriate retrieval information, and that if memory traces never changed, then it ought to be possible, in principle, to recover any memory, however old or weak, through appropriate retrieval cues. And these critics find such a state of affairs highly implausible. Again, the criticism is directed against a straw man. There is nothing in the hypothesis of cue-dependent forgetting that denies the possibility of permanent forgetting. It is quite conceivable that a particular retrieval environment necessary for the recovery of information from a given memory trace can never be reinstated or recreated, for a variety of reasons. Experiences remembered at one time and never remembered again can be explained as readily in the framework of cue-dependent forgetting as in the trace-dependent theories.

Then there is the claim that the hypothesis of cue-dependent forgetting constitutes the beginning of an infinitely regressive series and is therefore worthless. The statement that forgetting is a consequence of absence of appropriate retrieval information is regarded by certain critics as equivalent to the statement that forgetting of target events is caused by the "forgetting" of retrieval cues. And retrieval cues are forgotten because their retrieval cues are forgotten, and so on, ad infinitum.

The rebuttal of this criticism is relatively straightforward. The absence of retrieval information is no more equivalent to its being "forgotten" than its presence is to its being "remembered." Retrieval information plays the same role in remembering as, say, illumination does in the act of reading a printed page. Reading becomes impossible when the light is turned off, as remembering is impossible when relevant retrieval information is lacking. No one would claim that the absence of light, which causes failure of perception, is itself a consequence of the failure of perception. Only if one assumes that retrieval cues are remembered by the subject can one assume that they can also be forgotten. But the hypothesis of cue-dependent forgetting makes no such claim. A particular cognitive environment exists at the time of attempted retrieval of stored information; it exerts its effect on remembering independently of its origin, and an earlier cognitive environment can be restored within limits. It makes little sense, however, to think of this environment as being remembered or forgotten.

A more serious criticism of the hypothesis of cue-dependent forgetting is the complaint that it cannot be refuted. The validity of this criticism must be conceded, and it does detract from the attractiveness of the hypothesis. Nevertheless, the hypothesis is useful because it serves as a general principle that guides theoretical thinking, even if it is no substitute for more specific theoretical statements consistent with it that could be refuted by empirical data.

Another reasonably serious shortcoming of the hypothesis lies in the rather narrow domain of the experimental situations which have generated supporting evidence. The experiments presented above indicate something about the boundaries of this domain, although other types of positive evidence have also been reported (e.g. Blum et al. 1971; Warrington and Weiskrantz 1970). On the basis of the available evidence, it would be unreasonable to claim that the hypothesis of cue-dependent forgetting must be applicable to all kinds of forgetting. What is not yet clear is whether such a claim would be more unreasonable than the claim that, in situations in which reversal of forgetting through manipulation of the retrieval environment cannot be demonstrated, forgetting reflects some form of deterioration of memory traces.

Some critics of our hypothesis might want to argue that the reversal of forgetting demonstrated in single-trial free-recall tasks as well as in the multi-list retroaction experiment cannot be regarded as highly relevant. Such reversal, so the argument would go, only suggests that the presentation of category names as retrieval cues permits the subject to generate a restricted set of response alternatives upon which a recognition test can then be performed. This is the argument of the generation-recognition models of recall and recognition. The data from the generated-word recognition experiment (Tulving and Thomson 1973) make this particular criticism somewhat less convincing than it would have been only a few years ago. But even if one accepts the argument that reversal of forgetting "simply" demonstrates the superiority of recognition over recall, one still has to explain why recognition of an impoverished memory trace is possible while its recall is not. This brings us to the final criticism of the cue-dependent hypothesis.

This is the argument that the presentation of any retrieval cue that reverses forgetting has its effect through the trace rather than through the retrieval environment. The idea is that the retrieval cue somehow changes the trace, by updating the information it contains, by restoring some of the information lost from it, or, in some other manner, by "strengthening" it. McLeod, Williams, and Broadbent (1971), for instance, compared the effect of the second of two successively presented retrieval cues with that of the first cue in a list-item experiment. Since they found the second cue to be more effective than the first, it is possible to conclude that the presentation of the first cue "strengthened" the trace of the target word even when it was inadequate to bring about the recall of the word. However, since the second cue was presented under conditions in which the subject was aware of its relation to the first cue, a more parsimonious explanation of the results of the McLeod, Williams, and Broadbent experiment is that the second cue contained more information than the first one.

There is some evidence that is not easily reconciled with the hypothesis that retrieval cues serve the function of directly "strengthening" the stored traces. For instance, strong extra-experimental associates of list words as retrieval cues are known to facilitate recall of the list words (e.g. Thomson and Tulving 1970) but their presence interferes with recognition of list words (e.g. Tulving and Thomson 1971). When the subject is asked to remember a list that includes the word *butter,* he can recall the word better in the presence of the cue *bread* than in its absence. However, he can recognize *butter* better when it is given alone as a test item than when it appears in a test with the word *bread.* It is not clear how the hypothesis that cues strengthen traces would handle these experimental facts. Despite the existence of these and some other data that seem to be incompatible with the hypothesis that cues strengthen traces directly, the hypothesis does represent the central criticism of the view of forgetting as a cue-dependent process, and therefore it deserves more experimental and theoretical study and analysis.

The present paradigmatic view of remembering as a joint product of information from two separate sources might eventually be replaced by a different one. When that happens, the general problem of whether forgetting is a consequence of loss of trace information or absence of retrieval information, together with all its subordinate problems, might also have to be reformulated. But in the meantime, at least within the limited realm of memory for relatively unique events, the preferred view is of forgetting as an essentially cue-dependent phenomenon. The cue-dependent hypothesis is by no means free from objections, and many of its specific details remain to be worked out and tested in the laboratory. The objections are, however, no more serious than those that could be leveled against the alternative view of trace-dependent forgetting; and its advantages, including direct evidence in its favor, seem to outweigh the advantages of various trace-dependent theories.

References

Atkinson, R. C., and R. M. Shiffrin. 1968. Human memory: A proposed system and its control processes. In K. W. Spence and J. T. Spence, eds, *The Psychology of Learning and Motivation,* Vol. 2. N.Y.: Academic Press.

Bahrick, H. P. 1970. A two-phase model for prompted recall. *Psychological Review* 77:215–22.

Blum, G. S., J. R. Graef, L. S. Hauenstein, and F. T. Passini. 1971. Distinctive mental contexts in long-term memory. *International Journal of Clinical and Experimental Hypnosis* 19:117–33.

Freund, J. S., and B. J. Underwood. 1969. Storage and retrieval cues in free recall learning. *Journal of Experimental Psychology* 81:49–53.

Glanzer, M. 1972. Storage mechanisms in recall. In G. H. Bower, ed., *The Psychology of Learning and Motivation,* Vol. 5. N.Y.:Academic Press, pp. 129–93.

Keppel, G. 1968. Retroactive and proactive inhibition. In T. R. Dixon and D. L. Horton, eds., *Verbal Behavior and General Behavior Theory.* Englewood Cliffs, N.J.: Prentice-Hall, pp. 173–213.

Kintsch, W. 1970. Models for free recall and recognition. In D. A. Norman, ed., *Models of Human Memory.* N.Y.: Academic Press, pp. 331–73.

Light, L. L. 1972. Homonyms and synonyms as retrieval cues. *Journal of Experimental Psychology* 96:255–62.

McCormack, P. D. 1972. Recognition memory: How complex a retrieval system? *Canadian Journal of Psychology* 26:19–41.

McLeod, P. D., C. E. Williams, and D. E. Broadbent. 1971. Free recall with assistance from one and from two retrieval cues. *British Journal of Psychology* 62:59–65.

Postman, L., and G. Keppel. 1967. Retroactive inhibition in free recall. *Journal of Experimental Psychology* 74:203–11.

Postman, L., and K. Stark. 1969. Role of response availability of transfer and interference. *Journal of Experimental Psychology* 79:168–77.

Postman, L., and B. J. Underwood. 1973. Critical issues in interference theory. *Memory and Cognition* 1:19–40.

Slamecka, N. J. 1968. An examination of trace storage in free recall. *Journal of Experimental Psychology* 76:504–13.

Thomson, D. M., and E. Tulving. 1970. Associative encoding and retrieval: Weak and strong cues. *Journal of Experimental Psychology* 86:255–62.

Tulving, E., and Z. Pearlstone. 1966. Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior* 5:381–91.

Tulving, E., and J. Psotka. 1971. Retroactive inhibition in free recall: Inaccessibility of information available in the memory store. *Journal of Experimental Psychology* 87:1–8.

Tulving, E., and D. M. Thomson. 1971. Retrieval processes in recognition memory: Effects of associative context. *Journal of Experimental Psychology* 87:116–24.

Tulving, E., and D. M. Thomson. 1973. Encoding specificity and retrieval processes in episodic memory. *Psychological Review* 80:352–73.

Tulving, E., and M. J. Watkins. Continuity between recall and recognition. *American Journal of Psychology,* in press.

Warrington, E. K., and L. Weiskrantz. 1970. Amnesic syndrome: Consolidation or retrieval? *Nature* 228:628–30.

Waugh, N. C., and D. A. Norman. 1965. Primary memory. *Psychological Review* 72:89–104.

"Gottlieb, I think I know why we've been receiving so few commissions."