Psychological Review
1975, Vol. 82, No. 4, 261–275

Structure of Memory Traces

Endel Tulving
University of Toronto, Toronto, Canada, and Yale University

Michael J. Watkins
Yale University

A new descriptive theory of memory traces is presented. The memory trace is (a) conceptualized as a collection of trace elements and (b) defined in terms of the relation between the conditions and the product of retrieval. The properties of a trace thus defined are quantitatively described by measuring the gross, common, and reduced valences of two (or more) retrieval cues. These valences are determined by successively probing the target event with each of the cues. The data yielded by the successive probes are then used to construct the matrix or structure of the trace by means of the *reduction method*. The logic of this method, and hence the general theory, is applicable to a large variety of to-be-remembered material. A demonstration experiment showed that the structure of the traces of to-be-remembered word-events is sensitive to the conditions of initial encoding, and that forgetting of these events, under conditions of output interference, consists in a distinctive change in the pattern of trace elements. Some potential criticisms of the theory are considered.

Memory trace refers to a change in the memory system that results from the perception and encoding of an event. It constitutes a necessary condition for the subsequent retrieval of information about the event. A long-standing problem in the study of memory concerns the nature and properties of memory traces.

In a typical laboratory situation a person sees or hears a familiar word, which he is asked to store in memory. The perception of the word-event produces a change in the system, that is, it creates a memory trace. How do we conceptualize the trace of the word-event? How do we define it? How do we describe it quantitatively? And how do we relate our conceptualization, definition, and description of the trace to other ideas about memory?

This article suggests some possible answers to these questions by proposing a simple descriptive theory of the memory trace. Traces are (a) conceptualized as structured aggregates of trace elements, (b) defined as relations between the conditions and the products of retrieval, and (c) quantitatively described in terms of valences (effectiveness) of retrieval cues and of the interrelations of cue valences. The theory is embodied in a set of experimental and analytic procedures for specifying the structure of a memory trace.

Our conceptualization of memory traces as aggregates of elements is in keeping with currently popular notions of concepts and words being represented in memory as bundles of attributes or features (e.g., Anisfeld & Knapp, 1968; Bower, 1967; Morton, 1969; Underwood, 1969). We discuss trace elements, rather than features and attributes, solely for convenience in terminology.

We will not attempt to relate our trace-element approach to the other major theoretical tradition in characterizing memory representation, namely the associative network approach (e.g., Anderson & Bower, 1973; Kintsch, 1974; Rumelhart, Lindsay, & Norman, 1972). Our objective here is not to present a comprehensive review and analysis of the accomplishments of the various theories of memory representation; rather, it is to describe and discuss one possible approach to the problem of specifying properties of memory traces.

This research was supported by National Science Foundation Grant GB-40208X, and National Research Council of Canada Grant A8632.

Michael J. Watkins is now at the University of Toronto.

Requests for reprints should be sent to Endel Tulving, Department of Psychology, University of Toronto, Toronto M5S 1A1, Ontario, Canada.

TRACES AS RELATIONS BETWEEN OBSERVABLES

Since memory traces cannot be observed directly, they must be described indirectly, that is, in terms of things that *can* be observed. Three classes of observable entities are always involved in remembering: (a) the input into the memory system, (b) the retrieval query or cue, and (c) the output from the system. The relation between the input and the output was for a long time considered the only source of information about the memory trace. For instance, the correlation between some variable *X* (e.g., frequency, recency, etc.) and the probability of recall or recognition has been regarded as evidence that trace strength is a function of variable *X*. Thus, a property of the trace, its strength, has been defined in terms of the input–output relation.

There are at least three difficulties with the idea of traces as input–output relations. First, it encourages the notion of the trace as a copy of what has been presented to the memory system. The copy theory, even if rendered more plausible by assuming that the trace is but an incomplete or impoverished counterpart of the original object or event, seems intuitively unreasonable.

Second, a definition of the trace as an input–output relation cannot cope with the fact that the output from the system varies not only with the input but also with the conditions under which retrieval is attempted. This difficulty is usually avoided by observing the consequences of input only under invariant retrieval conditions. Unfortunately, this strategy precludes the study of some of the most interesting aspects of memory.

The third weakness of the input–output relation as a basis for defining traces is perhaps the most important one: It precludes making psychologically meaningful and empirically verifiable statements about the relations between input conditions and the trace of an item. If the trace is defined as a relation between input and output, its characterization must include a description of input conditions. Therefore, the relation between input conditions and the trace is a matter of logic, rather than experiment.

What we would like to have is some way of defining traces that would (a) permit us to arrive at a picture of the memory of an input event that is not a copy of the event, (b) reconcile the idea of a fixed trace with the fact that its activation will depend on the retrieval environment, and (c) foster questions about the relation between input conditions and the trace of an item whose nominal identity is held constant.

These objectives can be achieved if memory traces are defined in terms of the relation between conditions and products of retrieval, or as "a set of question-output relations" (Tulving & Bower, 1974, p. 294). The basic ideas are simple. Rather than thinking of the trace only as an aftereffect of the initial registration of an event, we think of it as a "cognitive blueprint" that specifies the conditions under which the recollection of the event will occur. Then, by assuming the encoding specificity principle (Tulving & Thomson, 1973)—the principle that any retrieval cue is effective only to the extent that its informational contents match the information contained in the trace of the event—we can interpret the observed effectiveness of different retrieval cues as evidence about the properties of the trace. Thus the definition of the trace in terms of the relation between the queries directed at the system and the output from the system leads to the description of the trace properties in terms of the effectiveness of various retrieval cues in producing recall.

Such an approach abates some of the difficulties inherent in the definition of memory traces as input–output relations. The description of the trace can emerge as something other than a copy of the input, since effective retrieval cues may be qualitatively quite different from the original event. Also, there would be no logical paradox about the fixed traces having variable manifestations, since the logic of the approach includes the assumption that the output from the system depends on the retrieval information (the cue) as well as on the stored information (the trace). Finally, this approach encourages questions about relations between input conditions and the resulting traces of one and the same item.

The major part of this article consists of the elaboration and illustration of this approach. But we will first briefly consider two currently popular approaches to the problem of describing memory traces.

TRACES AS FEATURE BUNDLES: TWO EXTANT METHODS

Many different methods have been used by various investigators in attempts to get at the properties of memory traces conceptualized as feature bundles. In a recent review, Tulving and Bower (1974) discussed 11 methods, and interested readers may wish to look at that review for a fuller background for the present discussion. We will mention here only two of the methods—those with which ours might be most directly contrasted.

One is Wickens's (1970, 1972) method of attribute shifts, which involves giving three or four successive short-term recall tests with the Brown-Peterson technique, with the study items sharing some particular perceptual or conceptual feature. Recall is typically found to decrease rapidly across the successive tests. A "critical" test is then given, in which the study items do not share one or more attributes that the items of previous tests had in common. The critical test often yields a high level of recall, compared with that for a control test in which no attribute change takes place. Wickens has argued that such a shift effect indicates that the systematically varied features were encoded into memory.

Although Wickens's attribute-shift method has produced a lot of interesting data, it does have a number of problems. One is that it appears to have quite limited applicability. So far, it has been used only within the confines of the conventional Brown-Peterson paradigm and under rather specific conditions of retention and intertest intervals. It is not known whether the method could be extended to measure the composition of traces of other sorts of input or with intervals measured in minutes, hours, or years. Another difficulty has to do with the fact that it is sometimes necessary to provide specific information to the subject before the attribute-shift effect is obtained—the change in the attribute is not always sufficient (Gardiner, Craik, & Birtwistle, 1972). Thus we are faced with the difficulty of deciding exactly what kind of retrieval information should be provided to identify trace features. Still another shortcoming of the method lies in the fact that the investigation of each potential encoding feature involves the use of a special set of items, a set which, as Underwood (1972) has pointed out, could serve to "prime" the feature under investigation. It follows that we cannot be certain that the dimensions of encoding suggested in the attribute-shift data have any generality beyond the precise conditions of the paradigm. Finally, the attribute-shift method provides no easy solution to the problem of determining the extent to which any two types of attribute shifts reflect the operation of a single underlying process. This is an aspect of the problem of establishing the relative dependence of trace features, a problem that occurs in many extant methods of describing traces (see Tulving & Bower, 1974).

Wickens's attribute-shift method is based on the implicit definition of the trace as an input–output relation. Although the input conditions in the method are specified in the form of differences between sets of to-be-remembered materials, and the output is indexed by a derived measure involving a comparison between two conditions, the effects of an attribute shift, and more particularly the ranking of the relative effectiveness of several attribute shifts (e.g., Wickens, 1970, Fig. 7; 1972, Fig. 3), do provide a picture of the trace as a relation between input and output. And, as we have seen, such a formulation has serious limitations.

The other method with which the theory and technique to be developed in this article can be contrasted is the method of false positive recognition errors. It was first described by Underwood (1965) and has been used as a tool for identifying component features of memory traces by Anisfeld and Knapp (1968) as well as many others (e.g., Bach & Underwood, 1970; Elias & Perfetti, 1973; Kosslyn & Bower, 1974). A person studies a list of items (words, sentences, or the like), and later takes a recognition test. Beside copies of study items and unrelated distractors, the test includes items that share particular features with study items. Thus, if one of the study items was the word CHAIR, the test might include the semantically related word *table* and the acoustically related word *hare*. To the extent that the person checks these distractor items as "old" more frequently than unrelated distractors, the features common to the study and the distractor items are implicated as a part of the trace of the study item. The features are usually designated by such terms as "associated words" or "rhyming words."

The method of false positive recognition errors is based on an implicit definition of the trace as a query–output relation. It is therefore free of the special difficulties inherent in input–output definitions of the trace. But it does have certain other difficulties. For instance, the experimenter gains information about trace features only to the extent that subjects make false positive "errors." Hence, the evidence for trace features depends, among other things, on the factors determining the frequency of such errors; yet some of these factors, such as the ratio of target to distractor items, probably have little to do with the nature of the trace. An attendant practical problem is that under typical conditions, where the subjects strive for accuracy, the density of the data that can be used to make inferences about trace properties is quite low, necessitating the testing of many subjects for stable results.

A second problem with the method is that the experimenter does not always know why the subject checks a particular distractor item as "old." If the study list contains the word CHAIR and the subject checks the distractor word *hare,* he may do so because of acoustic confusion with CHAIR, but he may also do it for some other reason, perhaps because of an idiosyncratically meaningful semantic confusion with some other study word. The very fact that control words in the test, selected to be unrelated to the study items, are checked as "old" suggests that the experimenter cannot always identify the basis for a false positive response.

A third problem has been discussed by Tulving and Bower (1974, p. 280) and will be only briefly mentioned here. Inferences about trace features require assumptions about features lost from the trace rather than features present in the test item. If the subject checks the distractor word *table* after having seen CHAIR in the list, he presumably does so because of an overlap of semantic features between *table* and CHAIR. But this overlap should only produce an "old" response to the extent that other features that normally distinguish *table* and CHAIR have been lost from the trace of CHAIR or are not extracted from the test stimulus *table*. It is not clear what kind of theory is needed here to relate the empirical data to the hypothetical memory trace.

Finally, there is again the problem concerning the relations among features, their independence and separability. Assume that the word CHAIR has been presented as a to-be-remembered item and that in the recognition test two distractors related to CHAIR, *table* and *sit,* are both checked as false positives. Does this mean that the trace of CHAIR includes one feature (e.g., "associative") or two features (e.g., "furniture" and "activity-instrument")? To rely on the experimenter's intuitions about the relation between two types of distractor items is not very satisfactory. Even where two items may appear to be very different (as with a semantically related word such as *table* and an acoustically related word such as *hare*), there is little to justify the conclusion that they have *nothing* in common. Evidence about the relations between different types of distractor items should be derived with an objective procedure, and so far no such procedure has been proposed.

RETRIEVAL CUING AND THE REDUCTION METHOD

Retrieval cuing as a method for describing memory traces has been briefly discussed by Tulving and Bower (1974). It has been used as a tool for gaining information about the composition of memory traces in a number of experiments (e.g., Barclay, Bransford, Franks, McCarrell, & Nitsch, 1974; Ghatala & Hurlbut, 1973; Light, 1972; Nelson & Brooks, 1974), primarily by providing a list of possible trace features or attributes. The present article, in essence, describes an extension of the method of retrieval cuing. The extension consists in the *successive* probing of a given trace with two or more different types of retrieval cues. By subjecting the results of such a procedure to a form of analysis that we call the *reduction method,* the relations among the retrieval cues with respect to the trace may be observed. And in the context of our theory, the observed cue relationships describe the structure of the referent trace.

Central to our approach is the encoding specificity principle mentioned earlier: A retrieval cue is effective to the extent that its informational contents match the informational contents of the trace. Thus we can infer informational properties of a particular trace and quantitatively specify its constituent trace elements by measuring the effectiveness of different retrieval cues.

At this point it is useful to introduce the idea of the valence of a cue, or more precisely the valence of a cue with respect to some particular trace. In fact we will define three types of valence: *gross valence* (often abbreviated to *valence*), *reduced valence,* and *common valence.* The gross valence of a cue with respect to the trace of a certain event refers simply to the probability with which that event can be recalled in the presence of the cue. In typical experimental situations where the subject's task is to reproduce the name of a word-event, gross valence refers to the probability of recall of the word in the presence of the cue. The valence will therefore assume a value between zero and unity. The valence of a cue must always be specified, at least implicitly, with reference to some defined trace, for the valence of a cue will vary widely for traces of different types of event—even among events involving the presentation of the same item under different conditions. Reduced valence refers to a cue valence that has been reduced by some other cue. The valence of cue X reduced by cue Y is given by the probability that the target can be recalled to cue X and cannot be recalled to cue Y. The common valence of two cues, such as cue X and cue Y, refers to the probability that the trace can be retrieved by cue X as well as by cue Y. It follows that the gross valence of cue X equals the sum of the reduced (by cue Y) valence of cue X and the common valence of cues X and Y.

Successive Probes

Consider now a simple experimental situation in which a subject is shown a familiar word among others in a to-be-remembered list. After the presentation of the list the subject is given two temporally separated cues for the target word in question, and asked to use the cues in recalling the target word. For instance, if the target word is CHAIR, the subject could be given the extralist cue *table,* perhaps with a statement of the conceptual relation of the cue to the target word. Thus, the complete cue might be "*table* (associated with) _____." Some time later in the test the subject is given another cue, such as "*hare* (rhymes with) _____." Of course, additional probes for the same target word could also be given, but we will consider only the situation in which two probes are used.

If these simple operations of presenting a word and probing it with two successive cues are replicated many times, frequencies of various possible outcomes can be used to estimate the valences of the cues. To do so, the data are summarized in the form of a 2 × 2 contingency table. One such data matrix from an imaginary experiment is shown in Table 1. Here one cue, X, was always presented as the first probe, and another cue, Y, as the second. The data in Table 1 tell us that cue X effected recall of the target item 45% of the time, and cue Y 60% of the time. We can also see that 40% of the time the target word was recalled to both cues, 5% of the time only to the first cue, and 20% of the time only to the second cue. Using our terminology, we can say that the gross valence of cue X was .45 and of cue Y was .60, that the common valence of cue X and cue Y was .40, and that the valence of cue X reduced by cue Y was .05, while the reduced valence of cue Y was .20.

In keeping with the encoding specificity principle, we regard the pattern of cue va-

TABLE 1

GROSS, COMMON, AND REDUCED VALENCES OF TWO CUES: HYPOTHETICAL DATA

First cue X	Second cue Y +	−	Total
+	.40	.05	.45
−	.20	.35	.55
Total	.60	.40	

lences shown in Table 1 as describing the trace of the target item. In other words, we have described the trace here in terms of the gross, common, and reduced valences of the two cues. Since the common and reduced valences are not predictable from the gross valences, the whole pattern of data tells us more about the memory trace than do the gross valences alone. Indeed, we might say that the pattern of the two valences tells us something about the *structure* of the memory trace.

To formalize the distinction between valences of cues and the properties of the trace, we will refer to components of the trace as elements and define them in terms of retrieval cues. Thus, Table 1 can be said to depict a trace containing 40% XY elements, 5% of $X\bar{Y}$ elements, 20% of $\bar{X}Y$ elements, and 35% $\bar{X}\bar{Y}$ elements. It is easy to imagine a situation in which the gross valences of cue X and cue Y are the same as those shown in Table 1, but the cell entries in the matrix are quite different. We would then say that even though the two traces had certain properties in common, their informational structures were different.

Reduction Method

Although the basic idea underlying the method of retrieval cuing for quantitatively describing memory traces is simple, its implementation is a bit more complicated. Complications arise because the valences of the two cues, X and Y, may depend on the order in which they are presented. This would be expected if the first cue changed the trace (retrieval-induced recoding of the trace), since the two cues would then be applied to different traces. If it were not for such potential confounding, we could let the data pattern in Table 1 represent the structure of the trace. But given that it is indeed possible, if not probable, that prompting recall of an event with one cue will change the valence of a second cue, another method for describing the trace is required. Our response to this requirement has been the development of the *reduction method.*

Consider again the basic experimental situation that we have described: A trace of a certain list item exists in memory and is probed successively with two cues, X and Y. But now the two cues are applied to the trace in two different cue-order conditions. In one condition the target is first probed with cue X and then, after other interpolated tests, with cue Y; in the other condition the order of the two cues is reversed. Data from the two cue-order conditions are tabulated separately in the form of two 2 × 2 tables. The data from these two tables are then combined into a third table, one that directly describes the trace structure.

Table 2 shows a set of data from an imaginary experiment in which gross, common, and reduced valences of two cues, X and Y, are obtained in two cue-order conditions as just described. The 2 × 2 table on the left (identical to Table 1) shows the results for the condition in which cue X was presented first and cue Y second, and the middle 2 × 2 table depicts the results with cue Y presented first and cue X second. We will refer to these tables as *data matrices.* The right-hand 2 × 2 table, which we will call the *trace matrix,* is composed of selected data from the two data matrices. As Table 2 shows, only two entries in each data matrix are used in constructing the trace matrix: the gross valence of the first cue and the reduced valence of the second. The remaining entries in the trace matrix are arrived at by subtraction, in keeping with the properties of 2 × 2 contingency tables.

The trace matrix is constructed in a way that minimizes the effect of cuing order and possible retrieval-induced recoding of the trace. Our basic assumption is that presenting a cue might cause a recoding of the target trace only if the cue is successful in effecting recall; a cue that does not elicit recall of the target item is assumed to leave the trace intact. Thus, in describing the

TABLE 2

TWO DATA MATRICES AND RESULTANT TRACE MATRIX: HYPOTHETICAL DATA

Data matrices								Trace matrix			
First cue X	Second cue Y			First cue Y	Second cue X			Cue X	Cue Y		
	+	−	Total		+	−	Total		+	−	Total
+	.40	.05	**.45**	+	.27	.25	**.52**	+	.32	.13	.45
−	**.20**	.35	.55	−	**.13**	.35	.48	−	.20	.35	.55
Total	.60	.40		Total	.40	.60		Total	.52	.48	

Note. Boldface indicates two entries from each data matrix (gross valence of first cue; reduced valence of second cue) used in constructing trace matrix.

trace we must take care not to use data that may be contaminated by the effects of a successful first cue. As Table 2 shows, no such contaminated data from the data matrices are used in constructing the trace matrix. The gross valence of the first cue is, of course, free of any contamination by the act of retrieval, and so is the reduced valence of the second cue, since this is based on only those cases in which the first cue did not effect recall.[1]

Two things should be noted about this method of estimating the gross, common, and reduced valences of the two cues with respect to a given trace. First, since we wish to avoid the potential problem of a change in the structure of a trace as a consequence of the act of retrieval, and since we therefore must exclude the data based on the retrieval of items previously retrieved, it is necessary to derive a trace matrix from two separate data matrices, each corresponding to one of the two possible test sequences. Provided our assumption is correct, this procedure gets around the difficulty created by the possible effects of retrieval on the trace. The effects of retrieving a trace on one occasion on retrieving it on a subsequent occasion represent an interesting research problem, but whatever these effects are they are not relevant to the reduction method.

The second point has to do with the problem of combining the data from the two data matrices into a single trace matrix. In the illustration shown in Table 2, the critical values for the two data matrices were entered directly into the trace matrix. Such direct transfer of observed data is possible if and only if the $\bar{X}\bar{Y}$ cells (probabilities of recall of target to neither cue) in the two data matrices are identical, as they are in Table 2 (.35 in both cases). But even if our basic assumption—that the proportion of items recalled at all (and hence the proportion not recalled at all) is independent of cuing order—is valid, small measurement errors or "noise" effects are to be expected. The critical entries in the data matrices will therefore usually require aligning before they can be combined into a trace matrix.

The most appropriate method of adjusting the data is simple to carry out, but in order to describe it unequivocally it is necessary to adopt a somewhat formal terminology. Consider the two data matrices and the corresponding trace matrix presented in Table 3. These matrices are shown in two forms, one in general symbolic notation, the other containing numerical data from a hypothetical experiment. The symbols X and Y refer to cues X and Y, and subscripts 1 and 2 indicate the order of presentation of the two cues. A plain cue symbol indicates successful recall of the target in the presence of the designated cue, while a symbol with an overbar indicates unsuccessful recall. Valence is in-

---

[1] The question has been raised whether the derived trace matrix represents the state of the trace at the time of the first probe or at the time of the second probe. There is no satisfactory answer to this question. Since the information in the memory store about any particular item probably changes in the interval between the two probes, our method can only be said to reveal the state of the trace as it "exists" at some time in the interval. The length of the interval between the two probes, of course, can be manipulated in experiment and the accompanying changes in the trace structure observed.

dicated by the symbol V. Thus, for instance, $V(Y_1)$ represents the gross valence of cue Y presented first, $V(\bar{X}_1Y_2)$ represents the reduced valence of cue Y, where cue X was presented first, and so forth. In the following discussion, the symbol V is omitted as redundant. In the trace matrix, the symbols X and Y have no subscripts, since they represent proportions of trace elements independently of cuing order. For instance, $\bar{X}$ represents the proportion of non-X elements in the trace, XY represents the proportion of XY elements, and so forth.

There is, of course, no need to adjust all the entries in the data matrices, but only those (critical) entries that are to be transferred to the trace matrix. The trace matrix entries are derived from the critical entries of the data matrices by using the following equations:

$$\bar{X}\bar{Y} = (\bar{X}_1\bar{Y}_2 + \bar{X}_2\bar{Y}_1)/2 \quad (1)$$
$$= (.38 + .32)/2 = .35,$$

$$X = X_1[(1 - \bar{X}\bar{Y})/(1 - \bar{X}_1\bar{Y}_2)] \quad (2)$$
$$= .43[(1 - .35)/(1 - .38)] = .45,$$

$$Y = Y_1[(1 - \bar{X}\bar{Y})/(1 - \bar{X}_2\bar{Y}_1)] \quad (3)$$
$$= .53[(1 - .35)/(1 - .32)] = .51,$$

$$\bar{X}Y = \bar{X}_1Y_2[(1 - \bar{X}\bar{Y})/(1 - \bar{X}_1\bar{Y}_2)] \quad (4)$$
$$= .19[(1 - .35)/(1 - .38)] = .20,$$

$$X\bar{Y} = X_2\bar{Y}_1[(1 - \bar{X}\bar{Y})/(1 - \bar{X}_2\bar{Y}_1)] \quad (5)$$
$$= .15[(1 - .35)/(1 - .32)] = .14.$$

The equations are illustrated with data from Table 3, with calculations rounded to two decimal places.

The logic underlying Equations 1 to 5 is as follows: The assumption that unsuccessful probing of the trace with a given cue leaves the trace unchanged implies that the order of administering the two cues does not affect the proportion of target items recalled to at least one of the two cues; that is, the commutative law is assumed. The assumption is stated formally in the following equation:

$$X_1Y_2 + X_1\bar{Y}_2 + \bar{X}_1Y_2 = X_2Y_1 + \bar{X}_2Y_1 + X_2\bar{Y}_1. \quad (6)$$

If Equation 6 is found to hold, then $\bar{X}_1\bar{Y}_2$ will equal $\bar{X}_2\bar{Y}_1$, and the critical values from the data matrices can be entered directly into the trace matrix. If Equation 6 does not hold, then it is assumed that the observed discrepancy between the left- and right-hand sides is attributable to sampling variability and measurement error. The best estimate of the true proportion of items recallable to at least one cue is then given by the mean of the two sides of Equation 6, and each side of the equation is adjusted to this best estimate. It is assumed that measurement and sampling errors are distributed proportionately among the terms of the equation, and hence the adjustment is made by multiplying each term by the ratio of the mean of the two sides to the value of the side of which the term is a part. This adjustment will render the equation true, and consequently produce two adjusted data matrices from which the critical entries can be used to form the trace matrix, exactly as was done in Table 2.

Rather than adjusting the data in the manner just indicated, in practice it is easier to work with the quantities $\bar{X}_1\bar{Y}_2$ and $\bar{X}_2\bar{Y}_1$, which are the complements of the values of the two sides of Equation 6. This is what has been done in Equations 1–5. The underlying logic is exactly as described, but the calculations indicated by Equations 1–5 are simpler.

Actually only three appropriately selected formulas given by Equations 1–5 need be applied; the other values can be obtained by simple subtraction, in keeping with the properties of 2×2 tables. In the illustrative data in Table 2, as well as in most actual experiments, the magnitude of adjustment resulting from the method just described is not very large, thus supporting our basic assumption and suggesting that the violation of the equality in Equation 6 is attributable to noise effects.

TABLE 3

DATA AND TRACE MATRICES

	Data matrices							Trace matrix			
First cue X	Second cue Y			First cue Y	Second cue X			Cue X	Cue Y		
	+	−	Total		+	−	Total		+	−	Total
Symbolic notation											
+	$V(X_1Y_2)$	$V(X_1\bar{Y}_2)$	$V(X_1)$	+	$V(X_2Y_1)$	$X(\bar{X}_2Y_1)$	$V(Y_1)$	+	XY	$X\bar{Y}$	X
−	$V(\bar{X}_1Y_2)$	$V(\bar{X}_1\bar{Y}_2)$	$V(\bar{X}_1)$	−	$V(X_2\bar{Y}_1)$	$V(\bar{X}_2\bar{Y}_1)$	$V(\bar{Y}_1)$	−	$\bar{X}Y$	$\bar{X}\bar{Y}$	$\bar{X}$
Total	$V(Y_2)$	$V(\bar{Y}_2)$		Total	$V(X_2)$	$V(\bar{X}_2)$		Total	Y	$\bar{Y}$	
Illustrative data											
+	.36	.07	.43	+	.32	.21	.53	+	.31	.14	.45
−	.19	.38	.57	−	.15	.32	.47	−	.20	.35	.55
Total	.55	.45		Total	.47	.53		Total	.51	.49	

EFFECTS OF ENCODING CONDITION AND RETENTION INTERVAL ON TRACE STRUCTURE: A DEMONSTRATION EXPERIMENT

We are now ready to explore the structure of traces created in experimental situations. We will describe a demonstration experiment in which encoding conditions and the length of retention interval were manipulated as independent variables and then analyze the cued recall data by the reduction method.

Method

Subjects were presented with six different lists of 16 to-be-remembered words each. Each to-be-remembered word was typed in capital letters and accompanied by a context word (input cue) typed in lower case letters. One half of the context words were associatively related to the to-be-remembered words, while the other half were words that rhymed with the to-be-remembered words. For instance, the associative (A) context for the to-be-remembered word CHAIR may have been *desk*, and the rhyming (R) context (presented with the same to-be-remembered word to another subject) may have been *pair*. The two types of context words defined the two encoding conditions of the experiment, and they appeared in the list in an essentially random order. Each list was presented visually just once, at the rate of 2 seconds per word pair. A short digit task was interpolated between the presentation and testing of each list.

Subjects were asked to study each word pair and to prepare for a cued-recall test, with the context words serving as cues for the to-be-recalled target words. In fact, to encourage subjects to continue to pay attention to the context words throughout the six lists, the target words were indeed tested with their respective input cues after the critical testing phase of each list; we will ignore the data from these input-cue recall tests, since they are of no relevance to our present concern.

The critical data were obtained from the extralist-cue tests interpolated between the presentation of each list and its test with the input cues. In these critical recall tests, each target word was probed twice, once with an associative extralist cue (e.g., *table*), and once with a rhyming extralist cue (e.g., *hare*). Subjects wrote their responses in a test booklet consisting of a number of pages. Different forms of the test booklet were used to manipulate relevant and control irrelevant variables. The first page of the booklet was used for recording the subject's name, the list number, and the digit string. Subsequent pages comprised the cued-recall test. With the exception of the last page of the booklet, which contained the input cues, each test page contained four rhyming cues and four associative cues. Each of the 16 targets was tested on two pages of the booklet, once with an associative extralist cue and once with a rhyming extralist cue.

Of the 16 targets, 4 were tested on the first two pages, 4 were tested on the first and third, 4 on the second and fourth, and 4 on the third and fourth pages. This manipulation of the position of retrieval cues defined the length of the retention interval, the second main independent variable in the experiment. Of each of these subsets of four targets, two had been encoded with respect to an associative context word and two with respect to a rhyming context word. One word in each of these pairs of targets was first probed with a rhyming cue and second with an associative cue, while for the other word the order of cues was reversed.

TABLE 4

TOTAL WORDS RECALLED TO ASSOCIATIVE (A) AND RHYMING (R) CUES AS A FUNCTION OF ENCODING FORMAT, CUING ORDER, AND PRESENCE OR ABSENCE OF FREE-RECALL TEST

Encoding format	Order of cues	Free-recall test	Both A and R (AR)	Only R ($\bar{A}R$)	Only A ($A\bar{R}$)	Neither A nor R ($\bar{A}\bar{R}$)
Associated word	A, then R	Yes	231	113	157	267
		No	242	96	159	271
	R, then A	Yes	270	62	172	264
		No	278	67	142	281
Rhyming word	A, then R	Yes	237	136	97	298
		No	232	126	97	313
	R, then A	Yes	234	104	121	309
		No	238	94	126	310

Note. The cell entries for each of the eight conditions sum to 768, for a total of 6144 subject-targets.

Thus for a given subject, each target word within a particular list represented a unique set of experimental conditions. The target words and conditions were balanced within each group of 16 subjects.

The nature of the relation of each cue to its intended target was explicitly specified in the test booklet: For instance, "*table* (associated with) ____," or "*hare* (rhymes with) ____." The relation between the two cues of a given target, however, was never specified. Thus, when the subject saw the cue "*hare* (rhymes with) ____" on page 2 of the booklet, he was not told that the target word sought here was, for instance, the same as the one that belonged to the cue "*table* (associated with) ____" on page 1.

An additional experimental variable concerned the presence or absence of a free-recall test before the extralist-cue recall test just described. A page containing free-recall instructions was interpolated between the first page and the extralist-cue test pages for half of the booklets of each subject.

A total of 64 Yale undergraduates of both sexes served as subjects. Since each subject was tested with six lists, and each list contained 16 target words, and each target was probed with two extralist cues, the critical data are based on 6,144 subject-words and 12,288 observations.

In summary, to-be-remembered words presented as members of experimental lists were encoded in the context of either associative or rhyming words. Each target word was subsequently probed with two extralist cues, an associative word and a rhyming word. The positions of the cues on the successive pages of the test booklet were systematically manipulated, and so was the presence or absence of an interpolated free-recall test. The data were treated by the reduction method.

Results

The results of main interest concern the effects of encoding conditions and retention intervals on memory traces obtained with the reduction method. As a general background to these results, Table 4 presents a summary of the total number of words recalled in the eight experimental conditions defined by the orthogonal combinations of the encoding format, the ordering of the associative (A) and rhyming (R) cues, and the presence or absence of the free-recall test; the data are pooled over retention intervals. The data for the free-recall conditions include all words, regardless of whether or not they were produced by the subject in the free-recall test.

Three points are worth noting. The first is that the free-recall test made little difference to the pattern of cuing results. For each combination of encoding format and order of cues, the data are practically indistinguishable for lists tested and not tested by free recall first. Thus, an initial free-recall test does not, or at least in this experiment did not, change the effectiveness of either the associative or rhyming cues.

A second point of interest in Table 4 concerns the effect of encoding conditions on the gross valences of the two kinds of cues. Encoding conditions affected the valence of associative cues, but not the valence of rhyming cues. The total number of words recalled to associative cues was 1,651 under associative encoding conditions and 1,382 under rhyming encoding conditions; the corresponding totals for the rhyming cues were 1,359 and 1,401. These data are very similar to those recently reported by Nelson, Wheeler, Borden, and Brooks (1974).

Finally, Table 4 shows that the violation of Equation 6 was negligible. Within each of the experimental conditions depicted in Table 4, the frequencies of responses in the $\bar{A}\bar{R}$ category (words recalled to neither A nor R cues) were quite similar for the two cuing orders.

For the analysis of principal interest, the data shown in Table 4 were entered into the data matrices, adjusted slightly in order to satisfy Equation 6, and then combined into trace matrices, one for each of eight experimental conditions. Table 5 lists for each of these conditions the four components of the trace, in the form of proportions of the relevant trace elements. The first four experimental conditions in Table 5 are defined in terms of the orthogonal combination of the free-recall and encoding format conditions. The trace components for these conditions were calculated on the basis of the data listed in Table 4. The other four conditions concern only those subject-lists for which the free-recall test was given; they are given by classifying the target items according to whether they had been produced in the free-recall test (some 32% had in fact been produced) on the one hand, and according to their encoding format on the other.

The data in Table 5 (a) confirm the implications in Table 4, both of the lack of any marked effect of the free-recall test and of the effects of the encoding format on the gross valence of the associative but not the rhyming cues, and (b) show that the memory traces of "more difficult" words, those not recalled in the free-recall test, are comparatively weak in the AR component of the trace.

Table 6 is concerned only with target words that subjects failed to produce in the free-recall test, and shows forgetting in the form of changes in trace structure during the course of the retention interval. The early testing refers to the targets cued on pages 1 and 2 of the test booklet, late testing to targets cued on pages 3 and 4.

Forgetting is seen to be reflected primarily in the reduction of the AR elements. Reduced valences of the two cues do not change very much during the testing period, and indeed the reduced valence of the rhyming cues (the $\bar{A}R$ component of the trace matrix) seems to have increased from the early to the late tests. The disintegration of the trace revealed by the decrease in the proportion of AR could reflect a transformation of AR elements into $A\bar{R}$ and $A\bar{R}$ elements at the same time that $A\bar{R}$ and $\bar{A}R$ elements are transformed into $\overline{AR}$ elements, but for the time being this is only conjecture.[2]

TABLE 5

TRACE MATRICES FOR TARGET WORDS IN VARIOUS EXPERIMENTAL CONDITIONS

Condition	Encoding format	Trace elements			
		$A\bar{R}$	AR	$\bar{A}R$	$\bar{A}\bar{R}$
No free-recall test given	Associative	.19	.33	.12	.36
	Rhyming	.16	.27	.16	.41
Free-recall test given					
All words	Associative	.22	.28	.15	.35
	Rhyming	.16	.26	.18	.40
Recalled words	Associative	.25	.55	.10	.10
	Rhyming	.15	.58	.16	.11
Unrecalled words	Associative	.21	.15	.17	.47
	Rhyming	.17	.11	.18	.54

Note. A refers to associative cues, R to rhyming cues.

EVALUATION OF THEORY

The theory and method for describing memory traces that we have outlined here, with gross, common, and reduced valences of two (or more) retrieval cues used to specify trace elements and hence trace structure, is really a rather simple-minded one. Indeed its simplicity might be regarded by some as its major weakness. The theory does not predict any phenomena, nor does it explain any, at least in the sense in which

---

[2] The fact that the AR elements of the trace are reduced in the course of forgetting and the $A\bar{R}$ or $\bar{A}R$ elements are not suggests that it is not useful to think of the items contributing data to the AR cell in the trace matrix as "easy" items or as having been recalled mostly by "good" subjects. Were such identification made, the problem would arise of how to explain the forgetting of "easy" but not "difficult" items, or the forgetting by "good" but not "poor" subjects. Any attempt to identify different cells in the trace matrix with items differing in difficulty or with subjects differing in ability also founders on the fact, inherent in the delayed test data in Table 6, that the conditional probability of recall to cue X, given failure of recall to the other cue, Y, is *higher* than the simple probability of recall to cue X uncontaminated by any earlier recall attempt.

TABLE 6

PROPORTIONS OF TRACE ELEMENTS DEFINED BY ASSOCIATIVE (A) AND RHYMING (R) CUES FOR WORDS NOT RECALLED IN FREE-RECALL TEST

Condition	Elements $A\bar{R}$	AR	$\bar{A}R$	$\bar{A}\bar{R}$
Associative encoding				
Early test	.22	.26	.14	.38
Late test	.19	.07	.24	.50
Rhyming encoding				
Early test	.14	.20	.18	.48
Late test	.18	.02	.23	.57

many people think of explanation. It represents little more than an elaboration of the definition of the memory trace in terms of the relation between the conditions and product of retrieval.

On the other hand, the theory enjoys several advantages over others. First, perhaps because of its simplicity, it can be used to define and describe memory traces of all kinds of events, using all kinds of retrieval cues. We have been concerned here with names of list words as the to-be-remembered materials, but there are no reasons why other aspects of list words, larger units of verbal materials, or other kinds of materials altogether could not be subjected to the same logic and the same experimental treatment. For instance, the copy cues (cues nominally identical with target items) used in a recognition test or the "invisible" cues assumed to guide retrieval in free recall could be treated in the same way as the A and R cues in our demonstration experiment. Second, since we define the trace in terms of query–output relations, its description need not emerge in terms of selected features, or as an impoverished copy, of the input item. Third, the trace matrix is a formal expression of the fact that the recall of an event varies with retrieval conditions. Hence there is no logical paradox between a fixed effect of a given experience and its variable memorial manifestations. Fourth, the trace matrix tells us not only about the relative efficiency of different types of retrieval cues but also about how they are related to one another. Finally, as we have shown in the demonstration experiment, it is a simple matter to find out how encoding conditions affect the trace of one and the same item. We hold the input items nominally constant, vary the conditions under which they are encoded, and then use the reduction method to arrive at the description of the trace.

Our theory as presented here is by no means complete. It will be revised and modified as we learn more about it and as we meet new problems in exploring its utility. In this final section of the article, we will consider some of the more obvious criticisms of the theory.

Trace Structure or Cue Effectiveness?

One of the questions that needs to be considered concerns the identification of various kinds of retrieval cue valences with properties of the memory trace and more specifically with the identification of our trace matrix with trace structure. Why do we refer to a pattern of cue valences as describing the structure of the memory trace, rather than simply discuss the data as they "really are," namely, as measures of cue effectiveness?

We invoke the concept of the memory trace because it helps to integrate, organize, and make sense out of what otherwise would be an unconnected list of experimental procedures and observations. The output from the memory system depends on previous inputs: The system can produce a certain output after but not before a certain input. This means that the system after the input must be somehow different from the system before the input. We can, like many others before us, think of this difference as the memory trace, and this primitive idea helps make sense of the changes in the output from the system. But if we wish to go beyond the mere assertion that an input into the system creates a new memory trace, we have to be more precise about how the trace is to be conceptualized, defined, and described. The theory outlined here attempts to do just this. It provides (a) an internally consistent set of statements about what the memory trace is and what its properties are, and (b) a set of objectively specified rules by which experimental observations are translated into theoretical terms.

The theory neither demands nor assumes any kind of "real" existence of some "thing" stored in the memory system. Memory trace, as defined by the theory and as described by the reduction method, represents a hypothetical construct that serves to organize a potentially "large list of relations between different questions directed at the system and the output from the system" (Tulving & Bower, 1974, p. 294). When these questions take the form of specific retrieval cues, the trace becomes the *referent* of the cues. We argued earlier that it makes little sense to talk about *the* effectiveness of a particular cue or *the* relation between two particular cues. Whenever we state the valence of a cue it is necessary to stipulate its referent, and for us the referent is the trace of a certain event. Thus we talk of valence of cue X with respect to trace T.

The concept of trace structure, and its description in terms of two or more cues and their informational overlap, represents a straightforward elaboration of these ideas. The measurement and specification of just the gross valences of individual cues provides only partial information about the changes that have taken place in the system as a consequence of a particular input. Probing the memory store with two cues aimed at one and the same target event, and observing the gross, common, and reduced valences of these cues, provides a richer picture of these changes. The similarity or relatedness of any two cues, X and Y, thus need not be determined by experimenters' intuitions; they can be measured, with respect to a specified trace T, by the reduction method. The magnitude of the common valence of two cues in the trace matrix, relative to the reduced valences, provides an objective measure of the extent to which the two cues contain the same retrieval information.

It is for these reasons, then, that we prefer to discuss the memory trace and its structure, rather than simply describe the observed relations between inputs into and outputs from the memory system.

Memory Trace and Memory Organization

Memory trace as defined and described by our theory always represents input events of a particular class. For instance, when input events consist of the presentation of individual words to many people, the resulting trace structure will refer to the population from which the subject-words were sampled, under specified conditions of presentation and recall. It would be a fairly straightforward matter to use just one item and many subjects, or many items and just one subject, and to describe the trace of classes of items defined in this manner. But it is not possible to specify the structure of the trace of a particular event in some particular subject, just as it is not possible to specify the state of one particular atomic nucleus at some particular instant.

Another issue concerns the fact that traces may be stored in an organized fashion (e.g., Mandler, 1967). A critic might wonder whether available evidence for organization in memory invalidates our method, which entails probing for individual items. There is, in fact, no conflict here. The organization of items in memory will be reflected in the trace matrices of these items. A corollary of this fact is that our method provides an approach to the study of organization in memory. For instance, our demonstration experiment showed different trace structures according to whether the target items had been presented in the context of associative or rhyming words (Tables 5 and 6). Further, our method readily lends itself to the investigation of the effects of intralist context effects on trace structure in a free-recall situation. Thus the organization of a given item, or set of items, can be manipulated—for instance, by adding conceptually related items to the list—and the consequent changes in traces described by the reduction method.

Cue Effects Through Mediation?

We have already noted that our theory was developed in the context of the encoding specificity principle: A cue is effective to the extent that its informational contents match the informational contents of the representation of the experienced episode. It is this principle that permits the specification of trace elements in terms of retrieval cues. But the assumption of a direct match between cue and whatever it is that has been encoded from an earlier event may be questioned.

It could be argued that the effect of a given cue could be mediated through some other cue, and therefore the valences of cues of a designated class provide evidence of questionable relevance to the informational contents of the stored representation of the to-be-retrieved event. The point can be illustrated with our standard example. Assume that the cue "*hare* (rhymes with) _____" effects recall of the target word CHAIR. In describing the reduction method and its theoretical underpinning, we have assumed that the cue *hare* is effective because its informational content—its interpretation by the system in semantic memory—somehow appropriately matches the informational content of CHAIR in episodic memory. It is possible to argue, however, that what may have happened is that the cue *hare* was used in the system to retrieve from semantic memory a number of rhyming words, including the word "chair." Following the implicit retrieval of "chair," its memory location could have been examined for evidence of a particular list tag (or the presence of certain contextual information) and the target word CHAIR recalled because an appropriate tag was identified. Alternatively, the system could have extracted relevant semantic information out of the implicitly retrieved "chair" and matched it with corresponding semantic information about CHAIR in episodic memory. In either case, the effect of the rhyming cue would have been mediated by semantic information, and hence the observed retrieval may tell us nothing about the overlap between the rhyming cue's own unique informational contents and the informational contents of the stored list item.

The reduction method makes it possible to evaluate the mediation hypothesis. The results of our demonstration experiment strongly suggest that, at least in this experiment, the degree of mediation was at most rather small. If a rhyming cue always effects retrieval of the stored information about the target item by first eliciting a word with semantic contents matching those of the target word, as illustrated in our example above, then the valence of the rhyming cue with respect to the trace reduced by the semantic (associative) cue should be zero. In fact, the proportions of $\bar{A}R$ and $A\bar{R}$ elements in the trace matrices were quite large. A simple mediation hypothesis also founders on the forgetting data shown in Table 6, since it could not explain why the common valence of the A and R cues (AR) drastically decreased over time, while the reduced valences ($\bar{A}R$ and $A\bar{R}$) either stayed the same or even increased.

There is a more fundamental retort to the potential criticism that the effect of a nominal cue may arise through the mediation of some other implicitly generated cue: Whether or not mediation occurs is of little relevance to the logic of our theory. The theory is not inconsistent with the notion that the mechanism of retrieval of stimulus information entails mediation. In fact, the theoretical formulation of the retrieval mechanism, in whatever form, would find relatively few constraints in the theory of trace structure that has been outlined here.

REFERENCES

Anderson, J. R., & Bower, G. H. *Human associative memory.* Washington, D.C.: V. H. Winston, 1973.

Anisfeld, M., & Knapp, M. Association, synonymity, and directionality in false recognition. *Journal of Experimental Psychology,* 1968, *77,* 171–179.

Bach, M. J., & Underwood, B. J. Developmental changes in memory attributes. *Journal of Educational Psychology,* 1970, *61,* 292–296.

Barclay, J. R., Bransford, J. D., Franks, J. J., McCarrell, N. S., & Nitsch, K. Comprehension and semantic flexibility. *Journal of Verbal Learning and Verbal Behavior,* 1974, *13,* 471–481.

Bower, G. H. A multicomponent theory of the memory trace. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 1). New York: Academic Press, 1967.

Elias, C. S., & Perfetti, C. A. Encoding task and recognition memory: The importance of semantic encoding. *Journal of Experimental Psychology,* 1973, *99,* 151–156.

Gardiner, J. M., Craik, F. I. M., & Birtwistle, J. Retrieval cues and release from proactive inhibition. *Journal of Verbal Learning and Verbal Behavior,* 1972, *11,* 778–783.

Ghatala, E. S., & Hurlbut, N. L. Effectiveness of acoustic and conceptual retrieval cues in memory for words at two grade levels. *Journal of Educational Psychology,* 1973, *64,* 261–266.

Kintsch, W. *The representation of meaning in memory.* Hillsdale, N.J.: Erlbaum Associates, 1974.

Kosslyn, S. M., & Bower, G. H. The role of imagery in sentence memory: A developmental study. *Child Development,* 1974, *102,* 579–584.

Light, L. L. Homonyms and synonyms as retrieval cues. *Journal of Experimental Psychology,* 1972, *96,* 255–262.

Mandler, G. Organization and memory. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 1). New York: Academic Press, 1967.

Morton, J. Interaction of information in word recognition. *Psychological Review,* 1969, *76,* 165–178.

Nelson, D. L., & Brooks, D. H. Relative effectiveness of rhymes and synonyms as retrieval cues. *Journal of Experimental Psychology,* 1974, *102,* 503–507.

Nelson, D. L., Wheeler, J. W., Jr., Borden, R. C., & Brooks, D. M. Levels of processing and cuing: Sensory versus meaning features. *Journal of Experimental Psychology,* 1974, *103,* 971–977.

Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model of long-term memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory.* New York: Academic Press, 1972.

Tulving, E., & Bower, G. H. The logic of memory representations. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 8). New York: Academic Press, 1974.

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

Underwood, B. J. False recognition produced by implicit verbal responses. *Journal of Experimental Psychology,* 1965, *70,* 122–129.

Underwood, B. J. Attributes of memory. *Psychological Review,* 1969, *76,* 559–573.

Underwood, B. J. Are we overloading memory? In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

Wickens, D. D. Encoding categories of words: An empirical approach to meaning. *Psychological Review,* 1970, *77,* 1–15.

Wickens, D. D. Characteristics of word encoding. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

(Received September 10, 1974)