

A propositional theory of recognition memory*

JOHN R. ANDERSON

University of Michigan, Ann Arbor, Michigan 41804

and

GORDON H. BOWER

Stanford University, Stanford, California 94305

This paper modifies the Anderson and Bower (1972) theory of recognition memory for words. A propositional representation is outlined for the contextual information underlying word recognition. Logical arguments are offered for preferring this representation over the undifferentiated associative representation used earlier. The propositional representation is used to interpret effects of verbal context upon recognition memory. The implications of these context effects are considered for two-process models of recall and recognition.

In a previous paper (Anderson & Bower, 1972), we presented a model for recognition memory and described the role it played in the recall of word lists. We argued that most recognition memory experiments require the S to retrieve information about past contexts in which the test item appeared. The item was recognized if the S decided that one of the contexts retrieved by the word was the list under test. Similar conceptions of recognition memory have been advanced by Bower, Lesgold, and Tieman (1969), Mandler (1972), and Kintsch (1973). The 1972 paper proposed a specific representation for the contextual information, consisting of a network of interassociated concepts. The 1972 representation basically was an extension of the memory structure of FRAN, a computer simulation model of free recall (Anderson, 1972). In FRAN, the concepts of the network were essentially words, and the associations were not differentiated according to their semantic role. However, as argued in Anderson and Bower (1973), such a representation does not impose sufficient structure on the information to account for the structure existing in a S's memory reports. Anderson and Bower present an alternative model, HAM. In HAM a careful distinction is maintained between words and concepts. Also, information is structured into propositions that specify the semantic relations between concepts. This paper will illustrate how the HAM representation deals with recognition memory for words and how it can be used to explain recent findings of context effects in recognition memory.

The HAM theory is partially embodied in a computer simulation program which operates in a question-answering task domain. One can interact with

the program from a Teletype, assert facts to it in a subset of English, ask it questions about earlier facts asserted, and receive appropriate answers. While the existence of this program influences the language in which we describe the theory, the reader should be warned that the program does not yet have all the properties of the theory. Program implementation has lagged considerably behind the process of theory construction.

A proposition in HAM is an associative configuration of elements which (a) is structured according to certain rules of formation; (b) has a truth value (i.e., it has something about which it is meaningful to say *true* or *false*); and (c) is abstract. This last point about abstractness needs some emphasis. It implies, for instance, that propositions should not be identified with the sentences that express them. Although strings of English words are commonly used to denote propositions, we would claim that words are never involved in abstract propositions. Rather, the elements of the proposition are the concepts which are expressed by the words. These concepts, like the propositions to which they belong, are *abstract*; they cannot be directly presented, but only expressed by a symbol or a word. This distinction between words and concepts turns out to have important empirical consequences.

HAM's propositional representation has marked similarities to that proposed by Rumelhart, Lindsay, and Norman (1972) and lesser similarity to those proposed by Kintsch (1972), Quillian (1969), and Schank (1972). Although the model applies to memory for individual words, it has its most natural and obvious application to memory for sentential material. Therefore, the initial description of HAM will illustrate how HAM stores sentences and retrieves answers to questions. With this as background, we can then show how the processes underlying recognition memory for single words can be viewed as basically the same processes as those underlying storage and retrieval of sentences.

*Preparation of this paper was supported by Grant MH-24360 from the National Institute of Mental Health to J. Anderson. We wish to thank Lynne Reder for her comments on this manuscript.

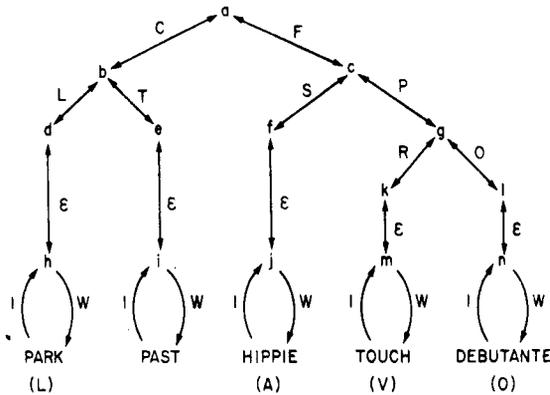


Fig. 1. A prototypical output of the parser. This structure will be encoded into long-term memory.

STORAGE OF SENTENCES

When a novel sentence is received, it is analyzed by a linguistic parser into a binary graph structure like that in Fig. 1. This graph structure is initially held in a working memory. A separate "tree" is composed for each proposition in the sentence. Because the example sentence, *In the park, the hippie touched the debutante*, involves only one atomic proposition, the graph structure in Fig. 1 is a simple tree. The uppercase letters label the arrows, while the lowercase letters representing nodes are arbitrary and simply facilitate reference to particular nodes. (Sometimes we represent nodes by unlabeled circles, however.) The labels on the associations indicate specific semantic relations holding among the nodes. There is a small well-defined set of these labels. Function words like *in* or *the* are not maintained in HAM's representation of the sentences. Any semantic information conveyed by choice of function words is expressed in the graph-structure configuration and by the choice of labels.

Each propositional tree is divided into two subtrees—a *context* subtree (arrow labeled C) and a *fact* subtree (labeled F). Intuitively, the nodes in the tree represent *ideas* and the links represent *relations* or *associations* between the ideas. So, in Fig. 1, node a represents the idea of the total proposition, node b the idea of the context,¹ and node c the idea of the fact. The proposition asserts that fact c is true in context b. The context node b is further divided into a *location* node d and a *time* node e, with arrows labeled L and T, respectively. Similarly, the fact node c is lead by an S arrow to a *subject* node f and by a P arrow to a *predicate* node g. That is, the fact is composed of a predication g being asserted about a subject f. Finally, in this example, the predicate node g is lead by an R arrow to the *relation* node k and by an O arrow to an object node l. So, what is being predicted of f is that it has relation k to l.

This completes the binary divisions in the tree. The node d represents a particular park, e a particular time in

the past, f a particular hippie, k a particular touching, and l a particular debutante. The general concepts of *park*, *past*, *hippie*, *touch*, and *debutante* are represented by the nodes h, i, j, m, and n, respectively. These general concept nodes are connected to the instance nodes by associations labeled with the set-membership relation, ε. These general concepts existed in the S's memory before receipt of the sentence and represent his idea of each concept. Connected to each node would be an associative structure (not shown) giving the meaning of the general idea. Note that these general ideas are distinct from the words *park*, *past*, *hippie*, *touch*, and *debutante*. The concepts are connected to the words by relations labeled W and the words to the concepts by relations labeled I.

All tree structure above the concept nodes is new and serves to record the novel information given in the input sentence. Thus, new propositions involve creating a complex configuration of known concepts according to certain labeled relations. To remember this sentence, each of the 13 working-memory links above the concept nodes must be acquired as associations in long-term memory.

Efficient access and use of stored information is as important a consideration in HAM as is the long-term representation of it. In HAM (see Anderson & Bower, 1973, Chap. 9), information is retrieved from memory in response to a probe, which is essentially a graphical specification of the pattern of information desired. Simple graph-searching techniques provide easy access to the information when it is directly stored in long-term memory. However, often the queried information is not directly stored, but it can be inferred from other propositions that are in long-term memory. Thus, when asked *Did Spinoza have a knee*, HAM would infer an affirmative answer from the facts *Spinoza is a man* and *all men have knees*. In Anderson and Bower (1973), we concentrated on fact retrieval and postponed extensive development of the inference system. Such inferential processes are a current theoretical concern. Much of the information retrieval to be discussed with respect to recognition memory is inferential in nature. For instance, upon seeing the test word *cat* the S may remember having seen that word printed in red ink and may know that all words in the study list were presented in red ink; he would thus infer that *cat* was one of the list words even though he did not directly remember this fact. This example involves probabilistic inference, whereas the previous *Spinoza* example involved pure logical deduction; but in either case, combinations of propositions are used to derive new facts. The inferential character of many recognition judgments has not been generally acknowledged in the literature. Instead, it has generally been assumed that recognition tests provide a sensitive, relatively uncontaminated measure of pure memory strength. The 1972 paper provided a rather extensive critique of strength theories, and there is no need to repeat the arguments here.

CONTEXTUAL PROPOSITIONS

Having outlined HAM's propositional system, the question can be answered of how it deals with the contextual information underlying list recognition judgments in list-learning experiments. In the 1972 theory, each fact about the context of presentation was recorded by a discrete element or a node in a memory network. These contextual elements recorded facts such as the color, location, and duration of the word's presentation, the S's internal physiological state, his mood, his thoughts, external stimuli, and any other feature to which the S attended during the word's presentation. Thus, the S might well store away something in memory equivalent to "While my stomach was gurgling and while fighting to keep my eyes open, I saw *dog* flashed on the screen in red capital letters. The word seemed to stay up for the longest while. I noticed the experimenter fiddling with the projector, and finally *cat* appeared on the screen." The problem with our earlier analysis is that such an autobiographical statement cannot be represented as an *unstructured* set of associations between elements or memory nodes. If one is to capture this complexity and the S's ability to reflect that complexity in his memory reports, one must resort to something like HAM's propositional representation.

Our basic claim will be that when a S hears a word, he propositionalizes ("describes") that event and stores one or more propositions in memory. If presented with the word *dog*, an obvious proposition that could be formed would be "In the List N context, I studied *dog*." Figure 2 illustrates how HAM might represent such a proposition. Connected to the root node *a* of that proposition are two other facts that the S has propositionalized while studying *dog*, namely, that his stomach was gurgling at the time and that he thought of a dog chasing a cat. Note that the proposition between *dog* and the other list word *cat* basically provides the S with the equivalent of an interword association.

An important aspect of this representation is that none of the propositions are connected directly to the word *dog* but rather to the concept node for *dog*. This reflects the assumption in HAM (see Anderson & Bower, 1973, Chap. 14, for an extensive discussion) that Ss tend to encode their memories about experiences with words in terms of the concepts expressed by the words rather than in terms of the words themselves.

When given a recognition memory test with *dog*, the S will use this propositional structure to decide whether the test word was presented on the study list. Given the complete structure in Fig. 2, the S can unambiguously conclude that *dog* is a list word because the structure records that *dog* was studied in the context of List N. However, some of the associations in Fig. 2 may not have been formed in the time allotted for study in a typical experiment (or may be unavailable at the time of testing). Therefore, the S may have to make a

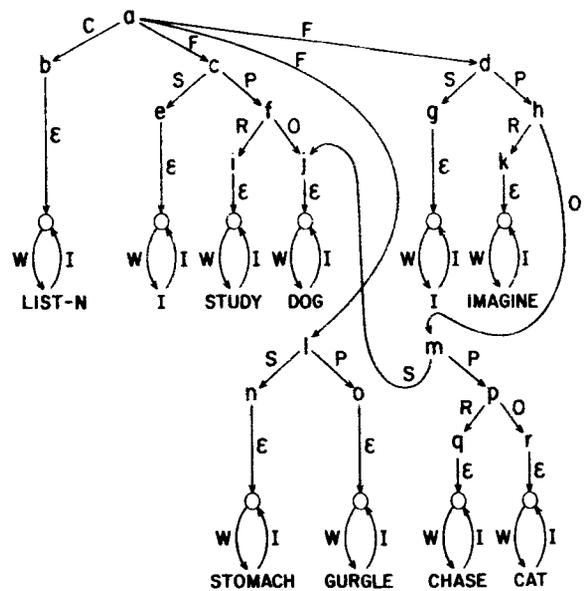


Fig. 2. Memory representation of propositions formed upon the appearance of the word *dog* in List N.

list-membership decision on the basis of an incomplete propositional structure in memory. For instance, the *e* association connecting node *b* (the context instance) to List N might not have been established. Then the S could no longer be certain that it was during List N that he studied *dog*. Perhaps the word *dog* occurred in some earlier experimental context. However, the S would still know that: (a) he had studied *dog* in some recent context, (b) during this time his stomach gurgled, and (c) he thought of a *dog* chasing a *cat*. Suppose the S also knew that his stomach was gurgling throughout List N or that *cat* was a List N word. (These facts would be recorded by other propositions in the S's memory.) Then the S could make the inference that *dog* occurred in List N with high confidence but not with the certainty he had in the former case.

The S might retrieve even less information about a word. For instance, he may only be able to remember that his stomach gurgled close in time to the appearance of the word *dog*. If only this information can be retrieved, naturally he will be much less confident that the word occurred in List N. In this manner, the S will display varying degrees of confidence in his recognition judgments. Such confidence ratings obviously show a positive relation to the probability that the word was actually in the list. That is, the S's degree of confidence reflects the amount of evidence he can find in memory for his decision.

The model outlined captures much of the phenomenological reports given by Ss in a list-learning experiment. Sometimes Ss will simply report that they "just know" the word occurred in the list. This would correspond to remembering the Fig. 2 proposition, *In the context of List N, I studied dog*. However, other times the Ss will give a more complex rationale of their

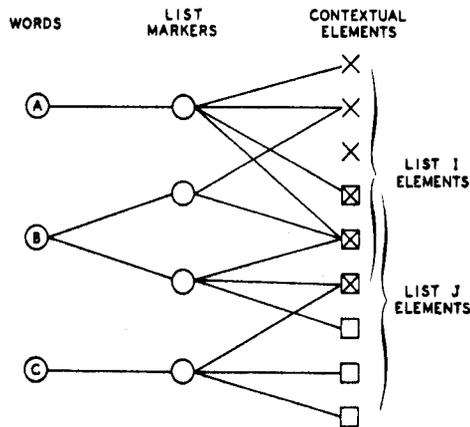


Fig. 3. A possible state of marking of three words, A, B, and C, with respect to List I and List J.

list-membership decisions which are not unlike our hypothetical example involving the association to *cat* and the stomach gurgle.²

There are three important aspects to the mnemonic representation illustrated in Fig. 2. First, there is the *word* itself. Second, there is the *dog instance*, *j*, which was created to encode the occurrence of the word in the list. In the 1972 paper, this instance was referred to as the *list marker*, a convention which will be continued here. If the same word occurred in another context, a different instance or list marker was created to record that occurrence. Thus, there is a one-to-one correspondence between each list marker and each occurrence of a word in a different context. From each list marker, the contextual propositions, which provide information relevant to a list discrimination judgment, can be accessed. These *contextual propositions* provide the third component in the mnemonic representation.

This conception can be reduced in essentials to the diagrams in Fig. 3, which is taken from the 1972 paper. In Fig. 3, nodes corresponding to words are connected to list markers, and list markers are connected in turn to samples of what were then called *contextual elements*. These contextual elements correspond to the contextual propositions of the new model. That is, they were intended to include such things as "physical characteristics of an item's presentation, implicit associations to the items, and some cognitive elements representing the list in question [Anderson & Bower, 1972, p. 101]." HAM's propositional representation makes explicit the information structure that was implicitly assigned to these contextual elements. Information like "stomach gurgling" or "I studied dog" are structures composed of elements and are not just unanalyzable nodes of memory. However, for analyses like those in the 1972 paper, the contextual propositions could be thought of as single elements.

Thus, for the level of analysis adopted in the 1972 paper, this new HAM representation is functionally isomorphic to the FRAN representation. The HAM

representation, however, permits a level of analysis that is much more complicated. The justification for this additional complication in the theoretical model is that it integrates a much wider variety of memory phenomena. It also remedies a number of empirical embarrassments to the old FRAN representation. We will now examine one such embarrassment.

ENCODING PROBLEMS IN RECOGNITION MEMORY

Several recent experiments demonstrate that even when the item has occurred in only one list context, access to information about its occurrence can be blocked by priming a different encoding of the word at the time of recognition testing. A landmark experiment demonstrating such encoding effects was performed by Light and Carter-Sobell (1970). They had Ss study homographs such as *jam*. An adjective modified these polysemous nouns to establish one or another interpretation of the noun. Thus, *jam* might be presented as *strawberry jam* or as *traffic jam*. The probability of recognizing the homograph was much higher if its study and test presentations used adjectives which suggested the same meaning of the noun. The obvious lesson is that it is not the physical word which is remembered in recognition memory experiments, but rather its meaning or sense. For other reasons we distinguish in HAM between the word as a physical stimulus and the various senses or concepts corresponding to that word (see Anderson & Bower, 1972, Chaps. 7 and 8). Each word was presumed to be connected to a *pushdown stack* of various senses and, conversely, each sense to a pushdown stack of lexical realizations. So, for instance, *ball* is at least two ways ambiguous—referring to a spherical object and a society dance. Similarly, the idea of a society dance may be expressed either by the lexical items *dance* or *ball*. The distinction between the word and the idea was emphasized with respect to Fig. 2. There we pointed out that HAM would tend to construct its contextual propositions about the word concepts and not about the words. Although not indicated in Fig. 2, each word could be connected to a number of senses and each sense to a number of lexical realizations.

The Light and Carter-Sobell experiment illustrates that semantic context can determine which sense or idea is activated and associated to a context tag. But let us consider a typical recognition experiment in which the word is presented alone for study and tested alone for recognition. It is reasonable to assume that the word sense that is highest on the pushdown stack of idea associations out of the word will be accessed and tagged at the time of study. At the time of recognition, this same sense will be retrieved and recognized. Evidence for this interpretation is provided by a recent experiment of Winograd and Conn (1971). They determined what were the relative frequencies of various interpretations or meanings for each of a set of homographs. All

homographs were studied without context, then were tested for recognition in one of three conditions. In a first condition, a sentence context was used that selected a high-frequency interpretation of the critical homograph (i.e., one likely to be on the top of HAM's pushdown stack); in a second condition, the contextual sentence selected a low-frequency meaning of the homograph; and in the third condition, the test homograph was presented in isolation. Recognition was equally high in the first and third conditions but lower in the second, where a low-frequency meaning had been selected by the test context. This result would be expected if the word sense likely to be selected when the word is presented in isolation is the high-frequency sense (i.e., in HAM, the idea association on top of the pushdown stack of associations from the word).

Some of the more striking results on recognition memory come from research that uses words that are not so transparently polysemous as the homographs used by Light and Carter-Sobell or Winograd and Conn. Such results have been reported by Tulving and Thomson (1971, 1973) and by DaPolito, Barker, and Wiant (1972). Tulving and Thomson presented to-be-remembered words along with strong or weak associations of the word. In the weak associate condition, Ss might study pairs like *train-black*, and in the strong associate condition, pairs like *white-black*. They found that subsequent recognition memory (for *black*) was better when the word had been studied and tested in the presence of the same associate (either high or low) rather than changing the cue word between study and test. In light of the results of Light and Carter-Sobell and of Winograd and Conn, the obvious explanation stems from the realization that even the word *black* has multiple senses. In the presence of *train*, one is likely to come up with the sense of *black* associated with soot and engine smoke or the sense associated with the glistening black of a polished toy train. However, in the presence of *white*, one is likely to come up with other senses of *black* such as a prototypical color or the absence of light or a race of people. This is essentially the interpretation Thomson and Tulving offer of their results: "The encoded engram of the unique event BLACK, in the context of 'train' and in the context of a specific set of TBR events, may be as different from the pattern of neural excitation corresponding to the generalized concept of BLACK as a beautiful and talented actress receiving an Oscar is different from any of the stars twinkling in the endless night [Thomson & Tulving, 1970, p. 261]."

It would seem that what distinguishes so-called homographs like *jam* from words like *black* is that the multiple senses of the former have relatively less semantic overlap. Nonetheless, the multiple senses of the "nonhomograph" are distinct, and if one sense is tagged during study, that sense must be retrieved later for successful recognition. Regarding this analysis, recent research by Reder, Anderson, and Bjork (in press) is

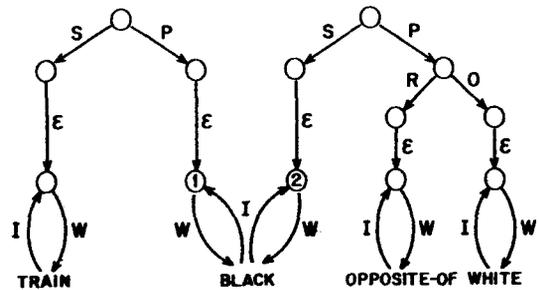


Fig. 4. Memory structure used in the disambiguation of *black*.

informative. High-frequency words like *black* which have multiple senses were contrasted with low-frequency words like *rhinoceros* which seem to have only one secure interpretation. In a partial replication of the Tulving and Thomson experiment, Reder et al found large detrimental effects for recognition and recall by switching associative context, but *only* for high-frequency words. The detrimental effects for low-frequency words were very much reduced and, in fact, were nonsignificant statistically. Thus, it would appear that the recognition difficulties reported by Tulving and Thomson are not due to change in context per se, but rather in context-induced changes in meaning.

An interesting question is how the current context activates a particular sense of a word. We suspect that one important mechanism may be an intersection search something like that employed for disambiguation in the TLC program of Quillian (1969). Such disambiguation techniques are not yet programmed in HAM, but it is clear how they would use the current context to select the appropriate meaning for a word. To illustrate, consider the fragment of memory structure in Fig. 4 showing two senses of *black*—one (node 1) involved in the proposition *The train is black* and the other (node 2) in the proposition *Black is the opposite of white*. For purposes of this example, let us assume that the former is not a particularly salient fact about *train* or about *black*. This low salience would be reflected in a low position on the pushdown stacks of associations leading between *train* and *black*. Hence, *black* would appear in the free association norms to *train* as only a weak associate. Nonetheless, let us assume that this proposition is the shortest, most available path connecting *train* and *black*. Now, suppose that *train* and *black* are presented simultaneously (as in the Thomson and Tulving experiment) and an intersection search is evoked. This search begins with parallel activation of all associative paths radiating out from *black* and *train*. This process concludes when the two sources of activation intersect. The point at which they intersect is a link on the shortest path between the two items. By assumption, the path between *black* and *train* in Fig. 4 provides this path of intersection. In selecting this path, the intersection search has selected a sense of *black* (namely, node 1) appropriate to the context of *train*. On the

other hand, the context *white* in Fig. 4 would select, by similar intersection techniques, the other sense of *black*, node 2.

This parallel, diffuse activation of all associations leading from a node is not adequate in general as a mechanism for searching long-term memory. In Anderson and Bower (1973, Chap. 12), empirical and logical arguments are advanced against the adequacy of this diffuse search. Similar logical considerations were stated earlier by Minsky (1968). The basic difficulty is that search of memory usually needs to be highly selective, considering only certain types of information. Nonetheless, one task for which the diffuse-activation technique seems useful is for the disambiguation of words in connected discourse. For instance, it explains how one immediately determines the appropriate sense of *bank* in the two sentences *He deposited his money in the bank* and *He was fishing from the bank*. An intersection search from *bank* to *money* or to *fishing* would appropriately interpret it as a monetary institution in the first case and as a location near a river or lake in the second. Lexical disambiguation was a prime concern in Quillian's original proposal for an intersection search. We would suggest that the context effects observed in recognition memory are produced by the same mechanism that produces lexical disambiguation in connected discourse. Thus, according to this analysis of encoding effects in recognition memory, the contextual phenomena with words is a manifestation of an extremely significant aspect of man's linguistic ability.

A Two-Process Model for Recognition?

In our earlier paper, we distinguished free recall from word recognition in that recognition was a fundamentally simpler process. That is, free recall involved (a) retrieval of the words by a search of long-term memory, and (b) recognition of the words by retrieval of contextual information. In contrast, we argued that recognition of complete units only involved (b). As evidence, we performed free recall experiments which dissociated these two components, increasing component (a) while at the same time decreasing component (b).

The existence of context effects in recognition memory (such as those reviewed) has recently been used to argue that recognition, too, must involve a retrieval process. Basically, the argument is that the one-process model of recognition makes the now untenable assumption that presentation of a stimulus provides automatic access to its trace in memory. As Tulving and Thomson (1971) state the matter: "they (findings of context effects in recognition memory) make it increasingly difficult, in our opinion, to cling to the view that there is something inherently different about processes of recall and recognition, and that an important aspect of the retrieval process, access to available information, is present in recall but not in

recognition [p. 123]." In light of evidence like Tulving and Thomson's experiment, one can hardly maintain that presentation of a word provides automatic access to the sense that was activated and tagged earlier. However, we would still maintain "the view that there is something inherently different about the process of recall and recognition." The reason is that it provides a facile way to discuss the many interactions between independent variables and mode of testing, recognition vs recall (see Anderson & Bower, 1972; Kintsch, 1970, 1973). While it is true that retrieval processes are centrally involved in either recognition or recall, the *types* of retrieval must be distinguished. To properly conceptualize the relation between recognition and recall, one must distinguish four different types of "retrieval": (a) the associative chaining through long-term memory during free recall, examining idea after idea, searching for senses of words that occurred in the list; (b) the examination of list markers or contextual propositions from a sense or idea in the attempt to determine whether that sense occurred in the list; (c) the generation of a lexical realization of the sense in recall; (d) the access to a sense from a word. Recall of a word involves retrieval aspects a, then b, and then c; on the other hand, word recognition involves retrieval d then b. So, although recognition and recall overlap in involving some common aspects of retrieval, they are distinct in terms of other retrieval subprocesses.

In our 1972 article, we assumed that both process c (access to the word from the sense) and process d (access to the sense from the word) would usually be successful and hence could be ignored. This is probably true when there are no special manipulations of concurrent contexts during study and testing and when single words are presented as whole units. However, it is well known that the sense-to-word transition can fail; this is clearly shown by the tip-of-the-tongue phenomenon (Brown & McNeil, 1966) or the "behind the mind's eye" phenomenon (Clayton & May, 1972). The research on context effects also indicates that presentation of the word does not guarantee access to its sense. These facts tell us that retrieval processes c and d can be made unreliable. However, we still feel that the hypothesis of infallible retrieval processes c and d is a useful *simplifying assumption* to make when analyzing free recall and recognition of words when they are not presented in special semantic contexts. The results of Winograd and Conn support this simplifying assumption for recognition memory.

In conclusion, we have argued that present evidence does not require rejection of the distinction between recall and recognition. On the contrary, the evidence suggests that further distinctions are required within the processes of recall and recognition.

REFERENCES

- Anderson, J. R. FRAN: A simulation model of free recall. In G. H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 5. New York: Academic Press, 1972.

- Anderson, J. R., & Bower, G. H. Recognition and retrieval processes in free recall. *Psychological Review*, 1972, 79, 97-123.
- Anderson, J. R., & Bower, G. H. *Human associative memory*. Washington, D.C.: Winston, 1973.
- Bower, G. H., Lesgold, A. M., & Tieman, D. G. Grouping operations in free recall. *Journal of Verbal Learning & Verbal Behavior*, 1969, 8, 481-493.
- Brown, R., & McNeil, D. The "tip of the tongue" phenomenon. *Journal of Verbal Learning & Verbal Behavior*, 1966, 5, 325-337.
- Clayton, K. N., & May, J. E. The behind-the-eye state. Paper read at the annual meeting of the Psychonomic Society, 1972.
- DaPolito, F., Barker, D., & Wiant, J. The effects of contextual changes on component recognition. *American Journal of Psychology*, 1972, 85, 431-440.
- Kintsch, W. Models for free recall and recognition. In D. A. Norman (Ed.), *Models of human memory*. New York: Academic Press, 1970. Pp. 307-373.
- Kintsch, W. Notes on the semantic structure of memory. In E. Tulving and W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press, 1972.
- Kintsch, W. Organization and retrieval in memory. Unpublished manuscript, University of Colorado, 1973.
- Light, L., & Carter-Sobell, L. Effects of changed semantic context on recognition memory. *Journal of Verbal Learning & Verbal Behavior*, 1970, 9, 1-11.
- Mandler, G. Organization and recognition. In E. Tulving and W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press, 1972.
- Minsky, M. (Ed.) *Semantic information processing*. Cambridge, Mass: M.I.T. Press, 1968.
- Quillian, M. R. The teachable language comprehender. *Communications of the Association for Computing Machinery*, 1969, 12, 459-476.
- Reder, L. M., Anderson, J. A., & Bjork, R. A. Semantic interpretation of encoding specificity. *Journal of Experimental Psychology*, in press.
- Rumelhart, D. E., Lindsay, P. H., & Norman, D. A. A process model for long-term memory. In E. Tulving and W. Donaldson (Eds.), *Organization of memory*. New York: Academic Press, 1972.
- Schank, R. C. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 1972, 3, 552-631.
- Thomson, D. M., & Tulving, E. Associative encoding and retrieval: Weak and strong cues. *Journal of Experimental Psychology*, 1970, 86, 255-262.
- Tulving, E., & Thomson, D. M. Retrieval processes in recognition memory: Effects of associative context. *Journal of Experimental Psychology*, 1971, 87, 175-184.
- Tulving, E., & Thomson, D. M. Encoding specificity and retrieval processes in episodic memory. *Psychological Review*, 1973, 80, 352-373.
- Winograd, E., & Conn, C. P. Evidence from recognition memory for specific encoding of unmodified homographs. *Journal of Verbal Learning & Verbal Behavior*, 1971, 10, 702-706.

NOTES

1. The reader should be careful to distinguish between our use of "context" in referring to information used in list discrimination and "context" in referring to this particular element of a propositional representation. The two senses of "context" do have important interrelationships, but they are by no means identical.

2. The fact that this theory corresponds to (and was in part suggested by) S introspections should not be held against it. Although some psychological phenomena have proven resistant to analysis through introspection, it should not be concluded that a S's introspections are never helpful. In any case, introspections constitute one set of data which the psychologist should try to explain. The complexity and open-endedness of S's memory reports argue for a propositional representation of the information used in word recognition. Simpler representations such as that provided in our 1972 paper are not capable of explaining these memory reports. A propositional representation is the only available system that is sufficient.

(Received for publication October 1, 1973;
revision received October 31, 1973.)