

Individuation and Reference in Memory: Proper Names and Definite Descriptions¹

JOHN ANDERSON
University of Michigan

AND

REID HASTIE
Harvard University

A theory is presented which assumes that individuals are represented by unique nodes in memory. To test the theory, simple facts were predicated of an individual person. Some facts referred to him by proper name, and other facts by his profession. In a *before* condition, subjects learned that the profession and name referred to the same individual before learning the facts, while in an *after* condition, they learned the identity after learning the facts. Subsequent to learning the facts and identities, subjects verified sentences based on what they had learned. Verification latencies indicated that in the before condition, one memory node was created to represent the individual, but two nodes were set up in the after condition. Asymmetries between proper names and professions indicate that the two types of referring expressions are treated differently in long-term memory.

Much of human knowledge is about "individuals"—that is, about specific times, places, people, objects, and so forth. This knowledge poses three questions for a theory of memory: How is this knowledge organized in memory? How is this knowledge used to make inferences about the individuals? How do linguistic descriptions and perceptual inputs contact the memory representations for individuals?

Traditional learning theories have ignored such questions. They have been concerned with how various stimulus properties (dimensions, attributes, features) come to control behavior. Recently there has been interest in how configurations of stimulus properties control behavior (e.g., Rescorla & Wagner, 1972). However, the question of how the organism comes to individuate—to confer on a stimulus configuration the status of an entity—has been ignored. Questions about individuals,

¹This research was supported by NIMH Grant No. MH-24360 to John Anderson, and by Yale Laboratory funds. We thank James Greeno and Lynne Reder for their comments on the manuscript.

properties and the relation between the two have been broached in recent computer simulation models of human memory (Anderson & Bower, 1973; Quillian, 1969; Rumelhart, Lindsay & Norman, 1972). Anderson and Bower (1973, Chapter 7) are quite explicit on these matters, dividing memory into nodes that represent individuals and nodes that represent concepts (or properties). The present paper reports an experiment designed to test some ideas incorporated in the Anderson and Bower model, HAM (*Human Associative Memory*).

Suppose HAM were presented with the following set of facts to learn:

- (1) James Bartlett is the lawyer.
- (2) James Bartlett rescued the kitten.
- (3) James Bartlett adopted the child.
- (4) The lawyer caused the accident.
- (5) The lawyer cursed the salesgirl.

(These facts are taken from material in the experiment we will report.) Sentence (1) indicates that the proper name, *James Bartlett*, refers to the same individual as the definite description, *the lawyer*. Sentences (2)–(5) assert four predicates about that individual. Sentences (2) and (3) use the proper name to refer to the individual and sentences (4) and (5) use the definite description.

Figure 1 provides a schematic representation of the knowledge structure that HAM would set up to encode this information. One node (labeled X) has been created to represent the individual under consideration. Attached to that node are the four predicates the subject has learned about the individual. Note that, although the facts in sentences (2) and (3) were learned via a proper name and in (4) and (5) via a definite description, they are all indiscriminately attached to node X. Thus we would expect subjects would not be able to recall which predicates were learned with which referring expressions. Anderson and Bower (1973, Chapter 9) found that subjects who learned facts like *George Washington*




FIG. 1. Representation of four facts about an individual.

had good health would say "yes" when asked if they had originally learned *The first president of the United States had good health*. They also tended to make similar "false alarm errors" when presented sentences with proper names, having actually heard sentences with definite descriptions.

The Anderson and Bower experiment was performed with American college students who certainly knew that such names and descriptions referred to the same individual. However, consider the case of someone who does not know the identity between the name and the description. He would regard facts about *George Washington* as facts about a different individual than facts about *the first president of the United States*. The present experiment is concerned with the question of how a subject would integrate the two distinct sets of facts when he learns that the proper name and the definite description refer to the same individual.

Philosophical analyses of the problem of reference have also considered the case in which a single individual is identified by two distinct labels. The classic illustration is the morning star-evening star example from the history of astronomy: The planet Venus was identified by two proper names, *the morning star* and *the evening star*. The discovery that the two "stars" were actually one planet was an astronomical finding of some importance. Frege (1892) used this illustration to argue against the position that proper names have no meaning other than their external referent. He proposed that if meaning were confined to the external referent, then the statement, *the morning star is the evening star* would be as trivial an identity as, *the morning star is the morning star*. Frege concluded that proper names must have meaning beyond their referent, and he advanced his famous distinction between the *reference* and the *sense* of an expression to solve this problem. HAM provides a solution to the same problem by its use of distinct individual nodes. In Frege's classic example, HAM would create two individual nodes, one for each proper name. Thus, although *the morning star* and *the evening star* have the same referent, their different meanings (or senses) are reflected by their distinct node structures in HAM.

In Fig. 1 we have assumed that the subject first learns sentence (1) and so creates one individual node connected to both *James Bartlett* and *lawyer*. However, suppose the subject did not learn this identity until after he had learned sentences (2)-(5). Then he would have to set up two individual nodes, a node X for facts about *James Bartlett* and another node Y for facts about *the lawyer*. The present research is concerned with how a subject would handle a sentence like (1) after creating two distinct individuals in memory.

In our experiment, subjects had to verify sentences such as *James Bartlett caused the accident*, which required them to make use of the identity between *X* and *Y*. Figure 2 shows two ways that HAM could represent the identity between *X* and *Y*. These two representations make different predictions about the times that it will take to make inferences. By an inference we mean that a predicate originally asserted of a definite description (e.g., the lawyer) is verified for a proper name or the reverse. In making predictions from the representations we will assume that reaction time increases with the number of facts that the subject must retrieve to verify a statement.

The solution presented in Fig. 2a is to encode, as a separate proposition, *X is equal to Y*. It is possible to infer *James Bartlett caused the accident* by chaining deductively through the three facts, *James Bartlett is X*, *X is equal to Y*, and *Y caused the accident*. In contrast, to verify *James Bartlett rescued the kitten*, the subject is only required to chain through two facts. Thus Fig. 2a predicts that inferences should always take longer than noninferences. In contrast, the representation in Fig. 2b predicts that inferences should take longer than noninferences only when the subject is asked to verify of a definite description a predicate originally learned about a proper name. In Fig. 2b the proper name has been attached to *Y*, the individual originally created for the description. The nodes *X* and *Y* can be regarded as referring to the same individual because they have the same proper name. Of course, it is possible that there are two *James Bartletts*, but as a rule, unless told otherwise, we regard two occurrences of individuals with the same proper name as two occurrences of the same

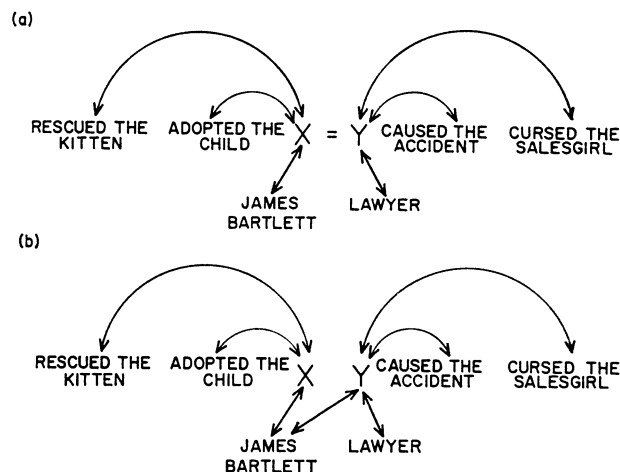


FIG. 2. Two possible representations of four facts about an individual in the identification-after condition.

individual. In Fig. 2b, it should take no longer to affirm of a proper name something that was originally asserted of the description than to affirm something originally asserted of the proper name. In fact, given HAM's mechanisms for organizing and searching long-term memory, we might expect that subjects would be faster with these inferences from proper name to description: In HAM it is a heuristic to search the more recent links first. Since the connection between *James Bartlett* and *Y* is more recent than the connection between *James Bartlett* and *X*, we would expect faster verification times for those predicates stored with *Y*. However, inferences from descriptions should take considerably longer. That is, to verify *the lawyer caused the accident*, the subject must execute the two-step sequence: *The lawyer is Y* and *Y caused the accident*; whereas, to verify *the lawyer rescued the kitten*, the subject must execute a four-step sequence: *The lawyer is Y*, *Y is James Bartlett*, *James Bartlett is X* and *X rescued the kitten*.

HAM also permits representations besides Figs. 2a and 2b. For instance, consider the opposite of Fig. 2b in which *lawyer* is connected to *X* as well as to *Y*. This representation is not acceptable because of the difference between definite descriptions and proper names. The relationship between a proper name and an individual node is one of reference but the relationship between a description and an individual node is one of set-membership. That is, *Y* is a member of the set of lawyers. Asserting that both *X* and *Y* are members of the same set does not imply that they are identical.

Yet another way to represent the identity of *X* and *Y* would be to combine the two into a single node, thus uniting the predicates attached to both *X* and *Y*. This would result in a representation such as the one depicted in Fig. 1, where there is no difference between inference and noninference conditions. Such a combination operation is completely foreign to HAM. Therefore, the failure to find a difference between inference and noninference would be a serious disconfirmation of HAM. A priori we have no basis to choose among Figs. 2a, 2b, and some of the alternates to HAM's representation. Our experiment was performed to gain empirical evidence to make this decision.

To summarize, this experiment involves three variables. The first, the *identification variable*, is concerned with when the subject is told about the identity of the proper name and the definite description: before or after learning facts about the individual. The second, the *inference variable*, is concerned with whether the predicate is originally learned with the label in the to-be-verified statement. The third, the *reference variable*, is concerned with whether the to-be-verified statement is asserted of the proper name or the description. These three variables

combine factorially to yield eight conditions. In the *identification-before condition*, we do not predict any difference between inference and non-inference sentences. However, a model that assumed the subject stored some relatively unanalyzed representation of each fact, specific to proper name or description, would expect inferences to take longer. In the *identification-after* condition, either of the proposed representations (Figs. 2a or 2b) would predict inferences would take longer for description labels, but it is an open question whether they should also be slower for proper names.

METHOD

Subjects

Sixteen subjects were recruited from the Yale undergraduate population. They were paid a wage of six dollars.

Materials and Design

The subjects in the experiment received information in the form of identification statements and predication statements about twelve artificial people. Each person had one identification statement which related his proper name and profession (definite description). These identifications were presented either before or after a set of 48 predications in the form of index cards containing both labels and a line-drawing portrait. The portrait was included to promote the formation of coherent, unitary impressions of the individuals and to emphasize the relation between the experiment and the natural acquaintance process. The twelve proper names were common male Anglo-Saxon names such as Robert Hawkins and James Bartlett. The twelve definite descriptions, randomly paired with the names, were common occupation titles such as lawyer and doctor. The 48 predications, four per artificial person, were made via five-word sentences containing a subject, verb, and object as in sentences (2)-(5). The complete set of proper names, definite descriptions, and predicates used in the experiment is in Appendix A. Two predicates were asserted of the individual's proper name, and two predicates were asserted of the description. Assignment of predicates to labels was random.

One-hundred-ninety-two test sentences were constructed from the study materials for the reaction-time verification test. These sentences were divided into two blocks of 96, so that within each block, there were 48 true sentences and 48 false sentences. The 48 true sentences in each condition were divided into the eight experimental conditions with six sentences in each. These conditions are defined by the orthogonal combination of the factors of identification, inference, and reference outlined

in the introduction. Because each condition was represented six times in each block, we could determine if the results held up over time. Thus, a fourth factor in the experimental design was delay, first block versus second block. No sentence in the two blocks was repeated. The 96 true sentences in the two blocks represented all possible combinations of the two labels with the four predicates for each of the twelve artificial people. Each block of 48 false sentences was created by randomly assigning each of the 48 predicates to different individuals. Each of the 24 labels (12 proper names and 12 professions) was used in two false statements in each block. Order of the sentences in each block was randomly determined. The materials and experimental conditions were counterbalanced in a number of ways: The order of the two blocks of 96 sentences was counterbalanced across subjects. The 12 proper name-description identification statements were separated arbitrarily into two subsets of six. The assignment of subsets to before and after conditions was counterbalanced across subjects. Response (true-false) assignment to telegraph keys was also counterbalanced across subjects. These three counterbalanced factors defined an eight cell between subjects design with one subject in each cell. Two completely independent sets of study and test materials were constructed from the same set of labels and predicates, and the eight-cell design was replicated with the second set.

Procedure

Subjects went through a three hour session which consisted of two principle parts: A learning phase during which subjects learned the material about the artificial people, followed by a test phase when the verification judgment task was performed.

During the learning phase, subjects were not informed that a memory test would follow. Instead, they were told that they were performing a sentence continuation task. They were given the set of 48 predications, one at a time, and required to "continue" each five word phrase in a natural manner. The instructions were written to suggest that the experimenter was interested in the psycholinguistic character of these continuations. This task led subjects to incidentally learn the information in the predications. It was adopted to prevent subjects from using unusual study strategies during the learning phase in response to real or imagined experimenter demand.

First, subjects were asked to memorize one set of six identifications. They studied a deck of six cards each of which contained a line drawing of a man's face, a proper name, and a profession. The experimenter drilled each subject until he was able to produce both name and profession to each picture on five passes through the deck. Then another drill

was performed in which the subject was given either a name or a profession and was required to produce the other label. This was repeated until the subject was able to respond ten times without error.

The subject was told that the set of to-be-continued sentences would refer to individuals whose names and professions he had learned, as well as to other individuals identified by name or profession but not both. He then performed the sentence continuation task with the set of 48 predications three times. Each time through the set, the order of the predications was randomized, and the subject was required to generate a new continuation of each sentence. Pilot studies indicated that three passes through the set would insure almost perfect performance on a recognition test. Following the continuation task, subjects learned a second set of six identifications under exactly the same procedures as the first set.

After learning the second set of identifications, the subjects proceeded to the verification phase of the experiment. Subjects were offered a bonus of one dollar if they could make fewer than 15% errors. Before starting the 192 test verifications, subjects ran through a series of practice verifications.

The procedure on each verification trial was as follows: Subjects sat approximately 2 ft in front of a screen. One second after a "ready" signal from the experimenter, the sentence was back-projected onto the screen. The slide was a white-on-black negative of a typed sentence, with the sentence spanning about 5 in. of the screen. The subject's right hand rested on a disc between two telegraph keys, and he responded true or false on each trial by pressing the key to the left or right. Four-inch diameter discs covered the target keys to insure that the subject would not miss. Reaction time was measured in centiseconds from the onset of the slide to the key press. The key press stopped the slide. If the subject's response was incorrect, the experimenter corrected him.

After 192 verifications, subjects were given a cued recall test of memory for the original predications. They were given a sheet containing a list of the proper names and professions, and were asked to write down, under each label, the two predicates originally learned to the label.

RESULTS AND DISCUSSION

The mean error rate for true sentences in the verification task was quite high (13.7%) and varied with experimental condition. This contrasts with a low (4.9%) error rate for false sentences. The HAM model would predict high error rates in the conditions which had long reaction times. In HAM (Anderson & Bower, 1973, Chapter 15), there is a cut-off time after which the memory search is terminated and a "false" response

is produced. Therefore, conditions with large reaction times will have high error rates. Thus we expected to observe similar trends in both error rate and reaction time data. The reaction times we will report are only for correct responses.

There were six observations for each of the sixteen conditions for each subject. These six observations were divided into sets of two observations from each of three artificial people. The mean of each subject's reaction time for these two-observation sets was computed. In a few cases (31 out of 768) both observations were errors and the mean was estimated from the means of the other two artificial people in that condition. These three means provided three reaction times for each combination of subject-by-experimental condition. These were the data used in the analysis of variance. Since assignment of materials to conditions was varied across subjects, the subject factor in the design also reflects differences in materials. In this way, we deal with the problem of generalization of our conclusions over materials (see Clark, 1973). This test is conservative in that subject-by-treatment interactions may include effects of the materials which are not present in overall means.

Table 1 presents the results of the analysis of variance on mean reaction times and error rates. Table 2 provides the mean reaction times and error rates for the sixteen conditions in the experimental design. The

TABLE 1
Analysis of Variance for Reaction Times and Error Rates^a

Effect	Reaction time		Error rate	
	<i>F</i> ratio	Significance	<i>F</i> ratio	Significance
D (delay)	2.97	.100	46.88	.001
IN (inference)	20.78	.001	13.39	.005
R (reference)	9.85	.010	1.84	.100
ID (identification)	3.65	.100	11.03	.005
D × IN	4.76	.050	18.81	.001
D × R	2.21	.100	1.65	.100
D × ID*	0.08	.500	10.75	.010
IN × R ₁	15.53	.005	80.69	.001
IN × ID	8.27	.010	0.88	.250
R × ID	0.07	.500	1.72	.100
D × IN × R*	0.43	.500	20.45	.001
D × IN × ID	0.28	.500	0.41	.500
D × R × ID*	0.51	.250	5.21	.050
IN × R × ID	22.53	.001	40.23	.001
D × IN × R × ID*	0.14	.500	10.60	.010

^a All tests of significance are based upon 1, 15 *df*.

TABLE 2
Mean Reaction Times (sec) and Error Rates

Block 1				
	Proper name		Description	
	Before	After	Before	After
Noninference	2.97	3.87	2.75	2.16
	10.4%	33.3%	4.2%	3.1%
Inference	3.35	3.70	2.92	4.10
	15.6%	9.4%	17.7%	57.3%

Block 2				
	Proper name		Description	
	Before	After	Before	After
Noninference	2.98	3.47	2.77	2.46
	7.3%	18.8%	9.4%	2.1%
Inference	2.87	2.76	2.57	3.87
	5.2%	2.1%	6.3%	16.7%

total complexity of the reaction time data will be addressed later by means of a mathematical model; at this point we will note only the most important effects. The overall standard error of the mean (based on the highest order interaction between treatments and subjects) was 0.28 sec for reaction times and 3.5% for error rates.

Figure 3 presents a partition of the data in terms of the delay, identification, and inference factors. In the before condition, there was almost no difference between inference and noninference verifications.




FIG. 3. The inference-by-identification interaction—block 1 vs block 2 of the test sentences.

In contrast, as predicted, subjects were considerably slower to make inferences in the after condition. Inspection of Table 2 indicates that this conclusion is limited to sentences with the definite description. The effects of the inference variable were reduced greatly from the first to the second block of verifications. There was a marginally significant decrease in the difference between inference and noninference verifications with delay, and a nonsignificant decrease in the magnitude of the identification-by-inference interaction.

A reduced interaction in the second half of the verification test is not surprising, since the inference effect in the after condition depended on the fact that the predicates were located at different individual nodes (see Fig. 2). After the first set of test sentences, the subject had experience with all 48 predicates, each occurring once in a false statement and once in a true statement. Thus, he would have had the opportunity to copy predicates from one individual node to the other. To the extent that such copying was occurring, the inference effect should have been diminished in the second half of the test.

Except for the significant delay by inference interaction, delay had no significant effects on reaction times. The significant three-way interaction between the identification, inference, and reference factors reflects exactly the pattern of results that would be expected if subjects used the Fig. 1 representation in the before condition and the Fig. 2b representation in the after-condition.

On the basis of Fig. 1, no differences were predicted between the inference and noninference conditions. A number of t tests were performed on differences between means from Table 2 (collapsed over the delay variable). These tests used as error terms the variance of subjects about the sixteen means in Table 2. In the before condition, the effect of the inference factor was not significant for definite descriptions ($t(225) = .04$) or for proper names ($t(225) = .47$). In contrast, the inference factor was significantly *slower* in the after condition for definite descriptions ($t(225) = 5.98$), and marginally *faster* in the after condition for proper names ($t(225) = 1.57$, $p < .10$). An analysis of the error rate data supports these conclusions: The effect of the inference variable in the before condition was not significant for definite descriptions ($t(225) = .44$) or for proper names ($t(225) = 1.49$). In the after condition, errors were much higher in the inference condition for descriptions ($t(225) = 9.83$) and lower for proper names ($t(225) = 5.80$).

Error rates reflected the reaction time results, so that all significant contrasts on reaction times in Table 1 were accompanied by error rate differences with the same sign. Some of the corresponding error rate effects were not significant, however. Error rates were much

more dramatically affected by the delay variable than reaction times. Error rates were reduced from 18.9 to 8.5% over the two blocks of verification sentences. This contrasts with a modest decrease in reaction times from a mean of 3.23–2.97 sec. Some of the significant error rate interactions with the delay variable did not reflect the pattern of means for the corresponding (nonsignificant) reaction time results. These discrepant interactions are starred in Table 1. The source of these inconsistencies is the large decrease across blocks of verification sentences for error rates in the description, after, inference cells of the design.

The correlation between error rates and reaction times in Table 2 is .77. If we correct for error in the measurement of reaction times, the estimated correlation coefficient is .88. Thus the error rates largely reflect reaction time as we predicted they would. However, even considering the corrected correlation coefficient, there was significant variability in the error rates not accounted for by a regression on reaction times ($F(14,225) = 4.05, p < .001$). The significant deviation from the linear regression is caused by the 57% error rate in the description, after, inference condition in block 1. When this point is removed from the analysis there is no significant deviation in a linear regression of error rates on reaction times.

The mean error rate in this one condition was over 50% which suggests that subjects were guessing. This would mean that correct responses are only "lucky guesses" not based on a successful search through memory, making the reaction time data in this condition meaningless. We argue that, on the contrary, correct reaction times reflect correct processing of information. We have no way to discriminate between these two possibilities in the present experiment. However, high error rates are frequently observed in tasks which require the subject to reason with a number of facts, and they are generally regarded as based on faulty reasoning and not random guessing (cf. Johnson, 1972).

The cued recall data provide further evidence that subjects stored facts about an individual at one node in the before condition, and at two separate nodes in the after condition. In the cued recall test at the end of the experimental session, subjects were given the labels (names and descriptions) and asked to recall which predicates had occurred with them in the initial learning task. If, in the before condition, subjects had stored facts indiscriminately at one node in memory, it should have been difficult to assign predicates to the correct labels. However, if in the after condition, subjects had initially stored the facts in distinct node structures according to labels, then recall by label should have been accurate. The cued recall data exhibit this pattern, with 16 out of 16 subjects making fewer label confusions for after identifications than for before

TABLE 3
Final Cued Recall Data (proportions)

	Before conditions	
	Cue for recall	
	Proper name	Definite description
Recall of a predicate to a proper name	.43	.31
Recall of a predicate to a definite description	.34	.51
Failure to recall or intrusion	.23	.19
	After conditions	
	Cue for recall	
	Proper name	Definite description
Recall of a predicate to a proper name	.69	.05
Recall of a predicate to a definite description	.05	.83
Failure to recall or intrusion	.26	.11

identifications (sign test, $p < .001$). The cued recall data are summarized in Table 3. It is clear that in the before condition there was considerable confusion in assignment of labels to predicates, while in the after condition, subjects were very accurate.

MATHEMATICAL MODEL

The reaction time data in Table 2 present a rather complex pattern of results. The analyses of variance reported in Table 1 indicate that many of the effects are highly significant. We claim that it implicates Fig. 1 as the memory representation for the before condition and Fig. 2b as the memory representation for the after condition. Figure 1 implies that there should be no differences between inference and noninference verifications. Figure 2b implies that inferences should take longer than noninferences only for verification sentences with the definite description label. For proper names in the after condition, because of the relative recency of the Y-proper name link, we predicted that reaction times

would be longer for noninferences than for inferences. While a model based on Figs. 1 and 2b captures many of the features of the reaction time data, it is not clear just how successful the model is. Therefore we will present a formal mathematical model, based on the figures and the process assumptions, and evaluate its success quantitatively.

The model requires five parameters. The first is p , the time to search a predicate. The second is w , the time to retrieve an individual from a proper name or a proper name from an individual. The third is d , the time to retrieve an individual from a definite description or the reverse. The fourth parameter is t , the time to encode the sentence, execute a true response, and perform other operations not associated with the memory search. The final parameter is a , which is a constant added to t for all times in the first block of verification test sentences. Thus we assume that there is a decrease in the encoding and response processes, but not in the memory search time, across blocks of test sentences.

Every sentence the subject is asked to verify specifies three elements: a subject, a verb; and an object. In HAM (see Anderson, 1974), simultaneous processes begin from each of these elements and independently search the memory structure. The time to verify a sentence is determined by the search process that first reaches a conclusion. Determining the first completion times for several racing processes can lead to complicated mathematical expressions. This certainly would be the case in a model for this experiment. An approximation to the predictions of this race model can be obtained by assuming that memory is searched from just one entry point, but that each entry point is equally likely to be the one from which the search is initiated. Since this approach yields a much more tractable model, we will pursue it here. In other words, we assume that memory is searched with one-third probabilities from subject, verb, and object.

Table 4 gives the model's predictions for the eight conditions of interest. We ignore the delay (block 1 vs block 2) effects accounted for by the a parameter in the model. Predictions for block 1 could be obtained by adding a to the equations, and predictions for block 2 may be taken directly from the equations. Individual conditions are referred to by the letter labels on the left-hand column of Table 4— B for before, A for after, P for proper name, D for definite description, N for noninference, and I for inference.

The first two rows in Table 4 represent inference and noninference verifications when the proper name is used in the before condition. Because it expects no difference between inference and noninference in the before condition HAM's predictions are identical for these two conditions. To explicate, consider the amount of search required from subject,

TABLE 4
 IIAM's Predictions for the True Data

Condition	Time for search originating from:				Mean prediction	Observed time	Predicted time
	Subject	Verb	Object				
1. BPN	<i>l + w + 2.5p</i>	<i>l + w + p</i>	<i>l + w + p</i>		<i>l + w + 1.5p</i>	2.97	3.20
2. BPI	<i>l + w + 2.5p</i>	<i>l + w + p</i>	<i>l + w + p</i>		<i>l + w + 1.5p</i>	3.11	3.20
3. BDN	<i>l + d + 2.5p</i>	<i>l + d + p</i>	<i>l + d + p</i>		<i>l + d + 1.5p</i>	2.76	2.67
4. BDI	<i>l + d + 2.5p</i>	<i>l + d + p</i>	<i>l + d + p</i>		<i>l + d + 1.5p</i>	2.74	2.67
5. APN	<i>l + 2w + 3.5p</i>	<i>l + w + p</i>	<i>l + w + p</i>		<i>l + 1.33w + 1.83p</i>	3.67	3.59
6. API	<i>l + w + 1.5p</i>	<i>l + w + p</i>	<i>l + w + p</i>		<i>l + w + 1.17p</i>	3.23	3.00
7. ADN	<i>l + d + 1.5p</i>	<i>l + d + p</i>	<i>l + d + p</i>		<i>l + d + 1.17p</i>	2.31	2.47
8. ADI	<i>l + d + 2w + 3.5p</i>	<i>l + d + 2w + p</i>	<i>l + d + 2w + p</i>		<i>l + d + 2w + 1.83p</i>	3.99	4.00

verb, and object in Fig. 1, which is the representation for the before condition. From the subject, the search will proceed from the proper name to X . This requires w seconds. From X , HAM would serially search the predicates, terminating the search when it finds the target predicate. Assuming that the target predicate is equally likely to be searched first, second, third, or fourth, there will be an average of 2.5 predicates searched before the correct one is uncovered. Adding these search components to the constant t , the expected search time from the subject is $t + w + 2.5p$. From the verb, HAM would search one predicate to get to X , and then from X , retrieve the proper name. We assume that the subject does not have to search through the other predicates attached to X to find the path from X to the proper name. In the HAM model, these predicate paths leading out of X could be ignored because they have a different logical relation than the desired path. Therefore the search from the verb will take $t + w + p$. The identical time is expected for a search from the object. The mean prediction given in Table 4 comes from a one-third weighting of each of these search times. The expected times and their derivations for rows 3 and 4 are identical to rows 1 and 2, except parameter d replaces the parameter w . This accounts for the time to proceed from an element like *lawyer* to the node X .

The derivations for the after conditions in rows 5-8 are somewhat more complicated than those for the before conditions. To understand these predictions, the reader should refer to Fig. 2b. The predictions in row 5 are for the condition in which a proper name is used and no inference is required. First, consider the predicted time for the search from the proper name: the link between the proper name and Y is more recently acquired, and hence should be more available than the link between the proper name and X . Therefore HAM's search will proceed from the proper name to Y , search the two predicates attached to this node, then back-up to X , and search one or two predicates at X before finding the target predicate. Thus two proper name-to-individual links must be searched as well as an average of 3.5 predicates, which gives the prediction in the subject column of row 5. From the verb or object, just one predicate will be searched to retrieve X and one link to retrieve the proper name. These search patterns give the predictions in the verb and object columns.

The predictions for the proper name in the after condition when an inference is required are in row 6. From the proper name, search proceeds to Y , and then one or two predicates must be searched to find the target. The verb and object estimates are the same as those in row 5. Row 7 is the definite description, noninference case. This is identical to

the proper name, inference case except that the subject goes between the profession and Y rather than between proper name and Y . Therefore, d replaces w in all predictions. Row 8 is the definite description, inference case. The search proceeds from the description to Y , and then the two predicates attached to Y are examined. Since the target predicate is not stored at Y , it is necessary to make an inference that X is Y , because they have the same name. This is the only situation where the subject must do more than search for a direct pattern match to the to-be-verified proposition. The inference from Y to X requires searching the link from Y to the proper name and the other link from the name to X . Once at X , one or two predicates must be searched to find the target. Thus, this search requires one transition from the description to Y , two transitions of the name-to-individual links, and search of a mean of 3.5 predicates. These components lead to the estimate under the subject column in row 8. From the verb or object, one predicate must be searched to retrieve X , two name-to-individual links must be crossed to retrieve Y , and the subject must proceed from Y to the description label.

The multiple regression estimates of the five parameters were: $w = 0.57$ sec, $p = 0.60$ sec, $d = 0.04$ sec, $a = 0.26$ sec, and $t. = 1.60$ sec. Table 4 gives the predicted times for each condition and the observed times averaged over block 1 and block 2. The model does a satisfactory job of predicting the verification times: With five parameters and 16 data points (there are 11 degrees-of-freedom in the model's prediction), the deviations of the model's predictions from the data are not significant ($F(11,225) = 1.06$).

The estimate of w is considerably larger than the estimate of d . The w parameter reflects the time to go between a word and an idea. Such a transition is also included in the transition from the profession label to the individual node and its estimated time is d . Perhaps the w parameter is larger than d because the link between the proper name and the individual is newly acquired, while the other word-to-idea connections are old and well-learned (e.g., between *lawyer* and the lawyer concept). Two factors contributed to the larger estimate of w . First, reaction times were considerably longer with proper names than with definite descriptions; the over all difference was 3.25 vs 2.95 sec. Second, reaction times were particularly long in those conditions (APN and ADI) where two name-to-individual links had to be traversed.

The model's goodness-of-fit is really quite remarkable. Still, we do not want to claim that the equations in Table 4 capture all the processes going on in the experiment. They only serve to show that a search model based on Figs. 1 and 2b does reproduce all the major features of the data.

Falsification Data

The false statements were not constructed to permit as detailed a partition of falsification times as we have presented for the true data. However, it is possible to partition these times according to two factors: first according to whether the label in the to-be-verified sentence was a proper name or a definite description and second according to whether the artificial person's identification was presented before or after the predication sentences during the learning phase of the experiment. As we might expect on the basis of our estimate of w (0.57 sec), subjects were much slower with proper names (3.39 sec) than with definite descriptions (2.86 sec). This effect was significant ($F(1,15) = 19.22$). There was also a marginally significant interaction between reference and identification ($F(1,15) = 5.59$, $p < .05$). Descriptions were slower in the after than in the before condition (3.03 vs 2.70 sec), and proper names were faster in the after than in the before condition (3.24 vs 3.52 sec). This interaction did not appear in the true data and would not be expected on the basis of our reasoning from Figs. 1 and 2b. We are at a loss to explain it.

CONCLUSION

The experiment provides support for HAM's representation of individuals in memory and its treatment of the proper name and definite description referring expressions. First, the lack of any difference between the inference and noninference sentences in the before condition is consistent with the claim that a single node is set up in memory to represent an individual, and that all facts about the individual are stored with this node. Second, the difference between the inference and noninference conditions in the after condition is consistent with the proposed discrete and disjoint character of a memory node. A subject cannot simply merge two nodes into one. Finally, the interactions with proper names and definite descriptions are consistent with the distinction made between their referential functions. The proper name was used to unify the X and Y nodes, but the definite description was not used in a similar way.

APPENDIX A

<i>Proper names</i>	<i>Definite descriptions</i>
John Anderson	the accountant
James Bartlett	the architect
Herbert Clark	the artist
Alan Collins	the banker

Robert Hawkins	the chemist
Neal Johnson	the dentist
Steven Madigan	the doctor
George Reynolds	the engineer
Edward Smith	the lawyer
Donald Thompson	the reporter
Richard Warren	the stockbroker
Michael Watkins	the teacher

Predicates

admired the dictator	invented the machine
adopted the child	kept his promise
assisted the invalid	kicked the dog
beat his son	littered the park
befriended the beggar	lost his temper
bored his companions	owned the pornography
bought the heroin	played the piano
broke the vase	protected the wildlife
caught the burglar	punished the secretary
caused the accident	received the prize
cheated the grocer	refused the bribe
comforted the victim	repaired the automobile
cursed the salesgirl	rescued the kitten
deceived his friends	resolved the dilemma
defended the blindman	ridiculed the cripple
discovered the danger	shoplifted the clothing
donated to charity	speaks four languages
entertained the guests	spread the rumor
failed the test	stole the wallet
forgave his enemy	supported his parents
helped the neighbor	teased the baby
hindered the fireman	threatened the woman
injured the pedestrian	won the game
insulted the tourist	wrote the bestseller

REFERENCES

- ANDERSON, J. R. Retrieval of propositional information from long-term memory. *Cognitive Psychology*, 1974, 6, 451-474.
- ANDERSON, J. R., & BOWER, G. H. *Human associative memory*. Washington, DC: Winston and Sons, 1973.
- CLARK, H. H. The language-as-fixed-effect-fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 1973, 12, 335-359.

- FREGE, G. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik*, 1892, 100. (Translation available in P. Greach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege*. Oxford: Basil Blackwell, 1960, pp. 56-78.)
- JOHNSON, D. M. *Systematic introduction to the psychology of thinking*. New York: Harper and Row, 1972.
- QUILLIAN, M. R. The teachable language comprehender. *Communications of the Association for Computing Machinery*, 1969, 12, 459-476.
- RESCORLA, R. A., & WAGNER, A. R. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory*. New York: Appleton-Century-Crofts, 1972.
- RUMELHART, D. E., LINDSAY, P. H., & NORMAN, D. A. A process model for long-term memory. In E. Tulving & W. Donaldson (Eds.), *Organization and memory*. New York: Academic Press, 1972.

(Accepted May 16, 1974)