

Memory for information about individuals

JOHN R. ANDERSON

Yale University, New Haven, Connecticut 06520

A semantic network model is structured so that usually there is only one node in the network to represent each individual. A series of experiments were performed to determine under what circumstances subjects would show unitary memory for individuals. The experiments were principally concerned with the speed with which subjects could retrieve the facts and make inferences from them. Subjects learned facts about individuals which could be referred to by two labels. The semantic network model predicted that subjects would integrate facts learned to one label with facts learned to the other. Evidence for such integration was found, but only when considerable effort was taken to encourage the subjects to develop a unitary impression of the individual. The situation was also investigated in which the subjects did not learn of the identity between the two labels until after the facts were learned to each label individually. There was evidence that subjects set up two nodes to represent the individual, one for each label. There was also evidence that, upon learning of the identity, subjects chose to abandon one of the two nodes and to start a process of copying information from the to-be-abandoned node to the preserved node.

A distinction can be made between concepts of individuals and concepts of general classes. In this discussion, "individuals" is used in a liberal sense to refer to specific times, places, objects, and so forth, as well as to specific people. "Classes" refers to variously defined groups of individuals. So, "Spot" would be an individual, whereas "dog" would be a class. The distinction between individuals and classes has long been important in metaphysical discussions (e.g., Frege, 1892; Loux, 1970; Ramsey, 1931; Russell, 1911-1912). The distinction between individuals and concepts is deeply entrenched in standard mathematical logic (e.g., Mendelson, 1964), where individuals are represented by constants and classes by predicates. The distinction has also been imported into recent psychological efforts to develop models of semantic memory (Anderson, 1976; Anderson & Bower, 1973; Collins & Quillian, 1972; Kintsch, 1974; Norman, Rumelhart, & the LNR research group, 1975).

Semantic network models such as those of Anderson and Bower and Norman and Rumelhart represent individuals differently than classes. There is one "type" node to correspond to each class. The class node may be connected to many "token" nodes, representing specific instances or individuals. In contrast, there are no special "subtoken" nodes connected to the individual

This research was supported by Grants GB 40298 and BMS76-00959 from the National Science Foundation. I would like to thank Rebecca Paulson for considerable assistance in the execution of each of these experiments. I would also like to thank Clayton Lewis for a suggestion that led to the design of Experiment 3. Finally, I would like to thank Lynne Reder for her many suggestions throughout this research program and for her comments on earlier versions of this manuscript. Correspondence concerning the paper should be sent to John Anderson, Department of Psychology, Yale University, New Haven, Connecticut 06520.

to represent different aspects of the individual. Thus, in a certain sense, an individual node is treated as an indivisible entity. The differential representation of individuals and classes is not logically necessary. As there can be different nodes for different instances of a class, so there might be different nodes for different appearances of an individual. However, these memory systems assume only one node per individual, reflecting the spatio-temporal continuity of that individual.

An experiment by Anderson and Hastie (1974) provided some evidence for this one individual-one node principle. They had subjects study sentences such as Sentences 1-5 given at the beginning of the experiment:

- (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.

Anderson and Hastie (1974) manipulated the point at which the identification statement (1) was learned relative to the predicate statements (2-5). In the "before" condition, subjects learned Sentence 1 before Sentences 2-5, while in the "after" condition, they learned Sentence 1 after Sentences 2-5. Subjects were then asked to verify sentences such as "James Bartlett rescued the kitten," which were directly stated, or sentences such as "James Bartlett caused the accident," which had to be inferred from the given facts. The mean verification times are displayed in Figure 1.

The data in Figure 1 are also classified by the label associated with the statement, either a proper name such as "James Bartlett" or a definite description such as "the lawyer." As can be seen in the before condition, there is very little effect on latencies of verifying inferences as opposed to direct statements. Anderson and Hastie (1974) proposed that subjects in the before




Figure 1. The verification data from Anderson and Hastie (1974).

condition set up a network representation illustrated schematically in Figure 2a. The proposal is that, upon hearing "James Bartlett is the lawyer," subjects set up an individual node "X" to which the proper name "James Bartlett" is attached and of which "the lawyer" is predicated. The predicates in Sentences 2-5 were attached to the individual node. The representation does not encode whether a predicate was learned via a proper name or a definite description. If subjects were verifying sentences from the data base in Figure 2a, no effect would be expected of inference vs direct statement, and none was observed in the before condition. This result would not be expected if subjects were setting up a distinct node for the proper name and definite description. Thus, the data support the one individual-one node hypothesis.

The data from the after condition in Figure 1 present a more complex pattern. There is relatively little effect of inference for proper names. If anything, subjects are slightly faster to make inferences. However, inferences to definite descriptions take almost twice as long as do those to direct statements. Figure 2b provides another schematic representation, similar to the one suggested by Anderson and Hastie (1974), which accounts for that pattern of data. Two individuals, X and Y, are set up to represent James Bartlett and the lawyer. To X are attached the predicates learned about James Bartlett and to Y, the predicates learned about the lawyer. Separate nodes are set up because the subject has no means of knowing that the terms are coreferential. When the subject is finally told that James Bartlett is the lawyer, he encodes this by the proposition "X turned out to be the same as (abbreviated TOTBSA) Y." The subject also introduces a link between the proper name and Y. Note that James Bartlett is directly connected to both X and Y. Therefore, there should be little difference between verifying a direct statement and an inference of the proper name. However, the lawyer is only connected directly to Y; it is connected to X only through the

proposition "X TOTBSA Y." Thus, to verify the inference "The lawyer adopted the child," the subject must reason "X adopted the child, X TOTBSA Y, Y is the lawyer." The long latencies to draw inferences for definite descriptions represent the need for this chain of reasoning.

While Figure 2b may account for the data, the reader may wonder whether it has any other motivation. Consider real world circumstances in which one might have to decide that one individual is the same as another. These are typically situations in which a new individual is described and only later does one realize the new individual is the same as an old individual. For instance, an acquaintance might describe some culprit who gave him a bad time. Upon this description, one might set up a new node to represent the culprit. Only later does one realize the culprit is his best friend, represented by another node in memory. Upon learning this, it seems reasonable to do two things. As in Figure 2b, one should record that the two individuals turned out to be the same. Second, it is unwise to maintain information about the same individual at two locations. Rather, it would be wiser to abandon one node, copying all information from it to the other node. That is, the memory system should try to reestablish the one individual-one node rule. This has been done partially in Figure 2b. Here, Node X has been chosen for abandoning and Node Y for preservation. The proper name has been copied to Node Y. The predicates from X have not yet been copied to Y, because they were not mentioned in the identification training involving Statement 1. There was some suggestion in the Anderson and Hastie (1974) experiment that the predicates were copied during the course of testing as they were encountered. The experiments here will report more decisive evidence on the issue of copying. Note in Figure 2b, that in copying the proper name link from X to Y, the old link involving X has not been deleted. This reflects an assumption that memory is not erasable.

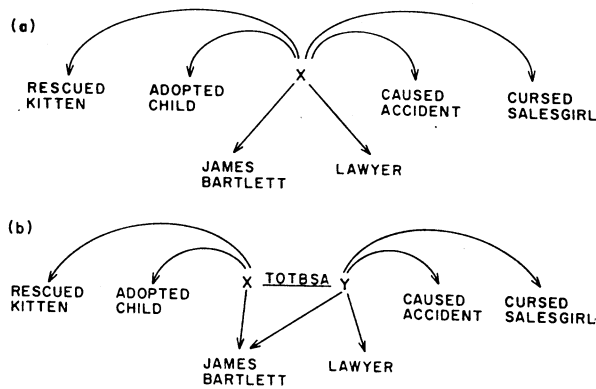


Figure 2. ACT representations for the before (a) and after (b) conditions of Anderson and Hastie (1974).

If this analysis of the before condition is correct, it becomes an interesting question which node is abandoned and why it is abandoned. In the earlier example, it seems that it would be better to abandon the culprit node and to maintain the friend node. The friend node probably has more information associated with it and would be a poor choice for abandoning. In Figure 2b, it seems reasonable to assume that X, the proper name code, is less meaningful and less elaborate than Y, the definite description code. The hypothesis of the present paper is that this is why X is abandoned and not Y; that is, nodes set up for proper names are less elaborated and, therefore, have less information to copy. One of the purposes of the experiments in this paper is to test this hypothesis.

The above analysis of the after condition is somewhat different than that offered by Anderson and Hastie (1974). They argue that the subject did not encode a proposition of the form "X TOTBSA Y," but rather that the two links out of James Bartlett were the mechanism used to encode the identity of X and Y. Anderson and Hastie's analysis seems unsatisfactory on two accounts. First, it seems bizarre to propose that subjects do not encode the identity between X and Y when they are explicitly drilled on that fact. Second, the links between the proper name and the individuals make a poor mechanism for recording identity. A proper name can be connected to multiple nodes with no implication of identity. The present author knows of no less than six John Andersons and yet suffers no confusions about their identities.

The dependent variable of principle interest in these experiments is reaction time. In interpreting the data, the ACT model is used. This model is described in detail and supporting evidence is given in Anderson (1976, Chapter 8). The model relates network representations such as those in Figure 2 to predictions about the time to retrieve information from them. It is assumed that, when a subject is tested with a probe such as "The lawyer adopted the child," activation is sent out from the concepts in the probe (i.e., lawyer, adopt, and child) and spreads through memory. When the total structure is activated that encodes the requisite information, the subject will respond "true." This is similar to the spreading activation model of Collins and Loftus (1975) and Collins and Quillian (1972). The structure that needs to be active in Figure 2b to answer this probe is "X adopted the child, X TOTBSA Y, and Y is the lawyer." The time to activate the structure will vary with two factors. First, there is the complexity or size of the structure. The more complex the structure, the longer it will take for the activation to spread and cover the structure. The second factor is the number of irrelevant paths leading away from the structure. Activation will spread down the irrelevant paths, dissipating the spread of activation along the relevant paths. For instance, all the irrelevant

paths leading from Y in Figure 2b would tend to dissipate the spread of activation from lawyer through Y to the predicate.

EXPERIMENT 1

The first experiment was an attempt to replicate the results of Anderson and Hastie (1974) with some important additions to the design of their experiment. Anderson and Hastie had taken pains to cause the subjects to create a unitary impression of a single individual in the before condition. In the initial phase of the experiment, where the subject learned about the identity between PN and DD, he was shown a sketch portrait of the individual to which the PN and DD labels applied. The present experiment deletes use of these labels to see if they were critical to the result in the before condition. It was hypothesized that, in the after condition, information attached to the definite description node is more richly represented than information attached to the proper name code. This hypothesis was tested by looking, before learning about the identity of the two labels, at verification of predicates to the proper name vs definite description. If definite description information is more elaborated, subjects should be more accurate in their judgment of this material.

The representations in Figure 2 make predictions about times to respond "false." Anderson and Hastie (1974) were not systematic in their gathering of "false" reaction times. This experiment will provide systematic analysis of the "false" data. Finally, this experiment doubled the length of the series of verification trials compared to that used by Anderson and Hastie. This provided more data relevant to the question of whether subjects are copying predicates from an abandoned node to a preserved node.

This is the first in a series of three experiments. Each of the three experiments involved a rather complex design, but the designs were all similar, so the details will be set forth only once.

Method

Subjects. Sixteen subjects were recruited from the general pool of subjects available at the Human Performance Center at the University of Michigan. They were paid \$2/h for their participation in the experiment.

Procedure. The experiment was run on the IBM 1800. The subjects sat before a CRT screen and responded to the various materials presented to them according to written instructions. There were six phases to the experiment. Before each phase, the subjects were given an appropriate set of instructions. The phases are illustrated with sample materials in Table 1 and are described below.

1. First identification. Subjects learned of the identity between four PN-DD pairs. They were drilled extensively on these identifications. The drill consisted of recalling the definite description (DD) label to a proper name (PN) probe and vice versa. The probe was presented on the screen and the subjects had to type in the correct response. They were required to

Table 1
Phases of the Experiment

1. First Identification	Learn: James Bartlett is the lawyer
2. Predicate Learning	Study: James Bartlett played the banjo James Bartlett drank the wine The lawyer repaired the car The lawyer sold the boat John Anderson built the house John Anderson grew a beard The doctor recycled the garbage The doctor watched the ballgame
3. First Reaction Time	Verify: James Bartlett played the banjo (true) The lawyer repaired the car (true) John Anderson built the house (true) The doctor recycled the garbage (true) James Bartlett built the house (false) The doctor repaired the car (false)
4. Second Identification	Learn: John Anderson is the doctor
5. Second Reaction Time	Verify: James Bartlett played the banjo (true) The lawyer drank the wine (true) John Anderson recycled the garbage (true) The doctor grew a beard (true) James Bartlett is the lawyer (true) The lawyer built a house (false) John Anderson repaired a car (false)
6. Final Recall	Recall the predicates originally learned to: James Bartlett The lawyer John Anderson The doctor

get all PN-DD pairings correct five times in a row before being permitted to pass to Phase 2 of the experiment. The four PN labels and the four DD labels that a subject encountered in this phase defined the material in the before condition.

2. Predicate learning. In this phase, subjects learned two predicates to each of nine PN labels and nine DD labels. Four of the PN and DD labels were the ones identified in Phase 1, four more of each were identified in Phase 4, and one of each were for practice purposes only. The predicates were simple verb-object combinations. Thus, the subject learned 36 sentences of the form "Subject verbed object," such as Sentences 2-5. Of the 36, 4 were intended for practice purposes. Order of all 36 sentences was randomized for each subject. The materials were learned in an incidental manner. The 36 sentences were presented one at a time on the screen and the subjects were required to continue each sentence in a manner consistent with its meaning. The subjects were not told that the experimenters were interested in the memory for the sentences, but rather that they were interested in the psycholinguistic character of the subjects' continuations. The incidental task was used to encourage the subjects to learn the material meaningfully. Subjects performed the continuation task for the set of 36 sentences three times.

The subject was told that the set of to-be-continued sentences would refer to individuals whose names and professions he had learned in Phase 1, as well as to other individuals identified by name or professions. He was told to think of the individuals as living in a small town in which they were known by profession or proper name or both.

3. First reaction time (RT). The subject was tested for his memory of the 36 sentences that he had learned. A test sentence

was presented on the screen. The subject pressed one of the two buttons, depending on whether he judged the sentence true or false. He was instructed to regard as true any sentence which followed from the material he had learned and as false all other sentences. The true sentences were the sentences he had studied, while the false sentences were created by re-pairing presented predicates with presented labels (not studied together). Care was taken not to create false sentences that were true inferences. That is, if PN and DD referred to the same individual, the subject was not presented a sentence consisting of PN as subject and a predicate learned to DD.

The subject first made eight verifications of true and false sentences based on the four practice sentences. Then he was tested with a sequence of 64 sentences, including the 32 experimental sentences as trues and 32 falses derived from these. The order of the 64 tests was randomized for each subject. One purpose of this phase was to guarantee that the subject knew the 32 basic sentences. If any errors were made on the 32 trues, subjects were given time to restudy these by generating another continuation. The error sentences were retested after the block of 64 sentences. In the retest, the error sentences were intermixed with randomly constructed falses. The subject did not leave the third phase until he had made one correct speeded response to each of the 32 sentences. A second purpose of this phase of the experiment was to gather evidence as to whether predicates were more poorly learned to proper names than to definite descriptions in the after condition.

The same procedure for gathering reaction times was used in all phases of the experiment. Subjects sat approximately 1.5 ft from the CRT screen. A sentence appeared on the screen spanning approximately 4 in. The subject's hands rested on two buttons, one designated as "true" and the other designated as "false." Subjects were permitted to choose the assignment of hands to buttons. The subject pressed one of the buttons to indicate his response. Reaction time was measured from the appearance of the sentence on the screen to the buttonpress and was recorded internally by the computer, along with the subject's response. With the buttonpress, the sentence disappeared from the screen and the feedback "You are correct" or "You are incorrect" appeared on the screen and stayed there for 1 sec. One second after its disappearance from the screen, the next sentence appeared for the subject's judgment.

4. Second identification. Subjects learned of the identity between four more pairs of PN and DD labels. These defined the materials for the after condition. The learning procedures were identical to Phase 1, involving the first identification.

5. Second reaction time. This phase collected measures of the subjects' speed and accuracy to judge statements based on the facts they had learned. The true material consisted of the 32 experimental sentences, 32 inferences obtained by pairing predicates learned with PN with the corresponding DD or vice versa, and 16 identification statements of the form "PN is DD" or "DD is PN." The identification statements were formed by testing both ways the four identifications learned in Phase 1 and the four identifications learned in Phase 4. Sixty-four false predicate statements were created by re-pairing labels with predicates and 16 false label statements were created by re-pairing proper names with definite descriptions. Altogether, there were 160 test sentences. These were preceded by eight practice sentences. The block of 160 test sentences was tested twice in two orders randomly determined for each subject. A short rest period intervened between the two passes. Subjects were given instructions that told them to anticipate inferences and identification sentences, and they were told to respond to these as true.

6. Final recall test. Subjects were shown the 16 experimental labels and asked to write the two predicates they had learned with each.

Materials. The study and test materials were randomly determined for each subject under the constraints of the design.

The materials were constructed from lists of 9 PNs, 9 DDs, and 36 predicates.

Results

Direct statements. One question of interest is what happened to the direct statements as the subject progressed through the experiment? These are the exact assertions he studied during predicate learning. Figure 3 provides the verification times and accuracy for these statements plotted as a function of experimental phase. The direct statements are plotted separately for definite descriptions and proper names (DD vs PN) for the before and after conditions (B vs A). The standard error of the "true" reaction times is 116 msec. This is computed from the Subject by Condition interaction and has 405 degrees of freedom. This is the error term that was used in computing all statistical tests. Because the materials were randomly generated for each subject, a test for generality over subjects also tested for generality over materials (see Clark, 1973).

The accuracy data show a general improvement over the course of the experiment. Mean accuracy was 79.9% in the first RT, 86.5% in the first pass through the second RT, and 91.8% in the second pass through the second RT. In part, this probably reflects subject adaptation to the task demands and, in part, real learning. Subjects were most accurate with DDA statements (94.9%), next with PNB (92.4%), next with DDB (89.8%), and least with PNA (78.3%). Thirteen of the 16 subjects performed better on definite description statements than on proper name statements in the after condition and one subject was tied ($p < .01$, sign test). In contrast, subjects were slightly worse in

their accuracy on definite description statements in the before condition. It seems intuitively reasonable that subjects should perform poorly to proper names in the after condition. A proper name without a profession is almost a nonsense syllable. It has less possibility for elaboration and establishment of a redundant memory trace (see Anderson, 1976, Chapter 10). The poorer performance to statements using proper names supports the hypothesis that subjects abandon the node of the proper name because it is less elaborated than the definite description node.

The reaction times show a general tendency to decrease over time: a mean of 2,317 msec in the first RT, 2,229 msec in the second RT, first pass, and 1,871 msec in the second RT, second pass. The linear trend was highly significant [$t(405) = 5.49$, $p < .001$]. Also, reaction times were longer to statements involving proper names than to those involving definite descriptions [2,292 msec vs 1,987 msec; $t(405) = 4.60$, $p < .001$]. Slower reaction times to statements involving proper names were also found by Anderson and Hastie (1974).

There was no significant overall difference between statements in the before and after conditions [2,152 msec vs 2,125 msec; $t(405) = .39$]. However, there were two interactions of this factor with time which are predicted from the ACT model (Anderson, 1976) and the representations in Figure 2. Reaction times were 328 msec faster in the after condition than in the before condition during the first RT, but an average of 126 msec slower during the second RT. This interaction was highly significant [$t(405) = 3.22$, $p < .001$]. This was expected, because, during the second identification (interspersed between first and second RT), the subject must add to his after representation structure, connecting the proper name and definite description nodes. The effect of this additional structure in the ACT model is to interfere with the retrieval of the original information attached to PN and DD. Note also PNA statements were more slowed (313 msec from first RT to second RT, first pass) than were DDA labels (10 msec). This difference was not significant by conventional standards [$t(405) = 1.32$, $p < .10$], but was replicated in Experiment 2 [$t(432) = 1.76$, $p < .05$]. The combined significance of the two experiments was quite high ($z = 2.17$, $p < .02$) (see Winer, 1971, p. 50). Greater interference for proper names than for definite descriptions in the after condition was predicted by Figure 2b. The interpolated training on the identity drilled in a link between the proper name and the individual node originally set up for the definite description. In the ACT model, activation spreading from the proper name would have to be split between the two paths. It is true that there are new interfering links from the definite description Node Y in Figure 2b. The links from the individual node should also slow down spread




Figure 3. Reaction time and accuracy to direct statement truths throughout the course of Experiment 1.

of activation. However, the interfering links are farther down the path from the definite description to its predicates than the interference on the path from the proper name to its predicates. It is a characteristic of a spreading activation model such as ACT that interference farther down the path will have less slowing effect than will interference earlier on the path. This is because the interference can be partially compensated for by the intersecting activation spreading from the predicate.

There was some tendency for the difference between before and after proper names to decrease from first to second pass in the second RT, although the effect was not significant [$t(405) = 1.01$]. This trend was also replicated in Experiment 2 [$t(432) = .75$]. Even the two experiments combined were not significant ($z = 1.25$). Experiment 3, which had a somewhat similar design, provided stronger evidence for the disappearance of differences between direct statements in the before and after conditions with repeated testing in the second RT. These data and other data that will be reviewed support the view that the representations in the before and after conditions become more similar with repeated testing.

Inferences and identifications. In addition to the direct statements, the second RT also involved tests of inferences and identifications. Table 2 provides a breakdown of the reaction time data for the true probes in the second RT. The data for the after condition in Table 2 nicely replicated Anderson and Hastie (1974). There was an inference effect only for definite descriptions. The before condition, however, did not replicate that of Anderson and Hastie. In the present experiment, there was a very strong inference effect for proper names [1.05 sec; $t(405) = 6.14$, $p < .001$] and a marginally significant effect for definite descriptions [$.31$ sec; $t(405) = 1.80$, $p < .10$]. The

difference between the two effects was significant [$t(405) = 3.07$, $p < .001$]. Some procedural change in this experiment must have been responsible for the large inference effect in the before condition. Later experiments in this paper help indicate which were the important changes.

There was a strong correlation (.83) between reaction time and error rate in Table 1. All major differences in reaction time involving inferences correspond to differences in errors. Therefore, the conclusions being drawn from reaction time are not compromised by possible speed-accuracy tradeoffs. Generally, throughout these experiments, reaction times and error rates show strong correlation.

It is of interest to note that the inference effects dissipated from the first to the second pass. The overall difference between inference and direct statement probes is .488 msec during Pass 1 and only .298 msec during Pass 2. This reduction was marginally significant [$t(405) = 1.65$, $p < .10$] and was replicated in Experiment 2 [$t(432) = .86$]. The experiments were significant in combination ($z = 1.77$, $p < .05$, one-tailed).

Analysis of "false" data. The false probes were constructed by combining a label with a predicate that it had not been studied with. The label can be classified according to whether it was proper name or definite description and according to whether it was identified in the before or after condition. Thus, there were four possibilities: PNB, PNA, DDB, and DDA. The predicate occurred in only one study sentence. The predicate can be classified as PNB, PNA, DDB, and DDA according to the label in its study sentence. False probes were constructed so that all 16 combinations of the two factors were tested equally frequently. Figure 4a presents the "false" data classified according to the label of the probe and Figure 4b presents the data classified according to the label associated with the predicate during study. The standard error of the reaction times in Figure 4 is 85 msec. There was a highly significant effect of the time at which false statements were tested. Average time to false statements increased 321 msec [$t(705) = 5.34$, $p < .001$] from first RT to second RT, Pass 1 and decreased from second RT, Pass 1 to second RT, Pass 2 by an average of 551 msec [$t(705) = 9.17$, $p < .001$]. The change in latencies for falses from first RT to second RT, Pass 1 was the opposite of the change for trues (88 msec; see Figure 3). This suggests that subjects became more cautious during the second RT in responding "false," because they then had to check for inferences.

The data in Figure 4a show that subjects were 375 msec slower [$t(705) = 7.64$, $p < .001$] to proper names than to definite descriptions. This is similar to the 305 msec difference for the "true" data in Figure 3. Subjects were also 123 msec [$t(705) = 2.51$, $p < .01$] slower to false statements involving after labels.

Table 2
True Reaction Times (in seconds) and Accuracy (in parentheses) for the Second Reaction Time of Experiment 1

	Before			After		
	DS	INF	IDENT	DS	INF	IDENT
	Pass 1					
PN	2.23 (.91)	3.28 (.78)	1.97 (.94)	2.62 (.70)	2.59 (.80)	1.91 (.95)
DD	2.05 (.90)	2.36 (.79)	1.89 (.98)	2.01 (.95)	2.63 (.67)	1.84 (.98)
	Pass 2					
PN	1.95 (.95)	2.18 (.90)	1.77 (1.00)	2.09 (.85)	2.28 (.91)	1.66 (.97)
DD	1.72 (.92)	1.95 (.88)	1.74 (.97)	1.73 (.95)	2.27 (.88)	1.68 (.98)

Note—In this table DA = direct statement, INF = inference, and IDENT = identification.




Figure 4. Reaction time to false probes throughout the course of Experiment 1: (a) classified by label of the probe; (b) classified by the label learned to the predicate of the probe.

The data revealed in Figure 4b, which gives classification by predicate, are somewhat more complicated than those in Figure 4a. In the before condition, there was little difference between predicates learned to proper names and predicates learned to definite descriptions. However, in the after condition, subjects are slower to predicates learned to proper name labels than to predicates learned to definite descriptions. Moreover, the size of this effect varied with time: 117 msec in first RT, 620 msec in second RT, first pass, and 185 msec in second RT, second pass. The overall difference was quite significant [$t(705) = 4.43$, $p < .001$]. The variation in the size of the differences was also significant [$F(2,705) = 5.15$, $p < .01$].

Some subjects reported that they decided probes were false by retrieving the label associated with the predicate. That is, when asked how they decided that it was false that James Bartlett wrote the formula, subjects reported that they would recall the label associated with the predicate "wrote the formula." This strategy was probably used because each predicate was uniquely associated with one individual, but each individual was associated with multiple predicates. As can be seen from Figure 2, there was a large

difference between PNA and DDA predicates in implementing this strategy. To decide that a probe involving a definite description predicate (e.g., "caused the accident") is false, it is only necessary to retrieve Y from the predicate and confirm that the label in the probe does not match PN or DD, which are connected to Y. However, the situation is not so simple for proper name predicates. The individual X connected to the proper name predicate is not directly connected to definite description. To decide such a predicate is false, it is necessary to retrieve X from the predicate, Y from X, and DD from Y. Thus, Figure 2 predicts, as is observed, that there should be a considerable difference in the second RT between falsifying probes with PNA predicates vs DDA predicates.

Final recall. The last task of the subject was to recall the two predicates he originally learned with each of the eight PN labels and with each of the eight DD labels. Table 3 summarizes this data. The data is classified by label, PN vs DD, and after vs before. The proportion of correct recalls of predicates, the proportion of confusions of predicates learned to the other label, and the proportion of intrusions or failures of recall are all listed. A confusion was only scored as such if the intruded predicate was the other label that had been identified with the cue label in first or second identification.

An interesting statistic is the ratio of confusions to correct recalls. This reflects the degree to which the subject was able to keep separate what predicates occurred with what labels. A ratio of 1.0 reflects total confusion and 0.0 reflects no confusion. The ratio was .52 in the before condition and .34 in the after condition. The difference between the two was statistically significant [$t(15) = 3.57$, $p < .001$]. It serves to indicate that the subject's memory representation at the end of the experiment better recorded the original sentences in the after condition. Such a difference would be expected if the subject's representation for the before condition was Figure 2a and for the after condition it was Figure 2b. Figure 2a does not record with which label a predicate was learned, while Figure 2b does. However, the difference between the confusion rates is rather small in absolute terms. To provide a comparison, in Anderson and Hastie (1974) the confusion rates were 68% in the before condition and 7% in the after condition. This supports the conclusion from the reaction time data that subjects

Table 3
Recall of Predicates in Experiment 1

	Correct	Confusion	Incorrect
Before			
PN	46%	25%	29%
DD	50%	25%	25%
After			
PN	53%	15%	32%
DD	53%	21%	26%

were not displaying much more tendency to integrate information in the before than in the after condition.

Discussion

Figure 5 presents schematically a set of information representations for the before and after conditions at various time points which will account for the available data. Consider the before condition: After the first identification in Phase 1, the subject has learned that proper name and definite description have the same referents. This is encoded by the links from DD and PN to the referent X. In the second phase, predicate learning, subjects attach predicates learned to the proper name to X. However, for some reason they choose to create a new referent, Y, for the definite description and to attach definite description predicates to this. This representation passed untouched through second identification and was the basis of verification in the initial part of the second RT. Note that DD is connected in this representation to both X and Y and to all predicates P₁, P₂, P₃, and P₄. Therefore, there should be little difference between an inference and a direct statement for definite description, and little difference was observed. In contrast, PN is not directly connected to Y, to which the DD predicates are attached. There is not even an indication of the equality of X and Y. Therefore, a large difference is predicted and was observed between inference and noninference. However, as a consequence of repeated testing, during the second RT, connections are built between X and Y. The

connections involving X come to be used exclusively and connections involving Y are weakened through disuse (indicated by dashes in Figure 5). Effectively, the subject decides to abandon Node Y for Node X. Therefore, toward the end of the second RT, the difference between inferences and direct statement begins to disappear.

The difference between this experiment and that of Anderson and Hastie (1974) is that subjects in the present experiment created a different individual node for the definite descriptions during predicate learning than the one they created during first identification. There was no necessary reason why they had to regard descriptions such as "the lawyer" in predicate learning as having the same referent as in first identification. It seems less likely that subjects would regard proper names such as "James Bartlett" as having different referents. Repetitions of professions often refer to different individuals, whereas repetitions of proper names less frequently do. This could explain the asymmetry in the inference effects for proper names vs definite descriptions (see Table 2). The marginally significant inference effects for DD may reflect those occasions when subjects set up a new node for PN as well as DD during predicate learning. The next experiment provides evidence that portrait drawings in the identification phase, which Anderson and Hastie used, do a lot to eliminate the inference effects in the before condition.

Figure 5 also illustrates the changes in information representation in the after condition. During predicate learning, two individuals, X and Y, are set up for the proper name and definite descriptions. To these individuals are attached the appropriate predicates. The effect of the second identification phase, as discussed earlier, is to initiate the abandoning of X, the PN node. An equality link is introduced between X and Y, and a link is introduced from PN to Y. The addition of these links leads to a slowing in the verification of direct statements in the second RT phase relative to the first RT phase. There is no inference effect for PN because PN is connected to all predicates. There is an inference effect for DD because it is not directly connected to all predicates. By the end of the second RT, the links involving X have been weakened (indicated by dotted lines) and new links have developed involving Y. This leads to the observed diminution of the inference effects.

An important feature about the information representations after the second RT phase is that both the before and after conditions preserve information about which predicate was learned to which label. However, the links that permit the distinction have begun to weaken. Therefore, some confusion is expected concerning the original assignment of predicates to labels. The data in Table 3 reveal a considerable amount of confusion in all conditions.




Figure 5. A schematic illustration of the change in information representation over the course of Experiment 1.

EXPERIMENT 2

Experiment 1 confirmed the analysis of the after condition of Anderson and Hastie (1974). It seems subjects choose to abandon the weaker proper name node. However, the inference effect in the before condition was unexpected and differs from the results of Anderson and Hastie. The post hoc explanation offered was that, in this experiment, subjects set up new nodes during predicate learning for definite descriptions. This may have resulted from the lack of portrait drawings during the first identification; the drawings may create a unitary impression. Therefore, a second experiment was performed that differed from the first only in the use of portrait drawings during first and second identifications.

Method

The experiment involved 17 subjects from the same pool as Experiment 1. The same procedures and materials were used with one exception: Phases 1 and 4, first and second identifications were split into two parts. The first phase involved presenting subjects with four portrait drawings and having them recall the proper name and profession to each. They were required to recall the name and profession 10 times to the drawing. Then subjects went on to the same drilling on the DD-PN identity that was used in Experiment 1.

Results

Since the results of this experiment were so similar to those of Experiment 1, they are presented without much discussion or reporting of statistical test. The reader may refer to Experiment 1 for discussion of the effects.




Figure 6. Reaction time and accuracy to direct statement trues throughout the course of Experiment 2. (Compare to Figure 3.)




Figure 7. Reaction time to false probes throughout the course of Experiment 2: (a) classified by the label of the probe; (b) classified by the label learned to the predicate of the probe. (Compare to Figure 4.)

The standard error of the reaction times to true probes was 104 msec ($df = 432$) and for the false probes it was 69 msec ($df = 752$). Figure 6 traces the performance to direct statement trues over the course of the experiment; it shows a similar pattern to Figure 3 for Experiment 1. Figure 7a provides an analysis of false probes classified by label and Figure 7b provides an analysis of false probes classified by predicate. The "false" data show a similar pattern to Figure 4.

Table 4 provides an analysis of reaction times to all true probes in the second RT. Note that the results of Anderson and Hastie (1974) for the after condition were again replicated: Subjects only showed inference effects for DD labels. An inference effect was also obtained for the before condition for both PN and DD labels. The average inference effect was .32 sec for the first pass, which was significant [$t(432) = 3.08, p < .001$]. The effect was .68 sec in Experiment 1, however. The difference between the two effects was quite significant [$t(837) = 2.32, p < .01$]. Similarly, the accuracy difference was much less in Experiment 2 [12% for Experiment 1 vs 3.5% for Experiment 2; $t(31) = 2.74, p < .01$]. So it seems that the use of portraits considerably reduced but did not eliminate the inference effect in the before condition.

Table 4
True Reaction Times (in seconds) and Accuracy (in parentheses) for the Second Reaction Time of Experiment 2

	Before			After		
	DS	INF	IDENT	DS	INF	IDENT
	Pass 1					
PN	2.28 (.81)	2.53 (.85)	1.93 (.97)	2.58 (.69)	2.49 (.88)	1.86 (.99)
DD	1.89 (.90)	2.28 (.79)	1.71 (.94)	1.86 (.90)	2.77 (.51)	1.62 (.99)
	Pass 2					
PN	1.69 (.90)	1.94 (.88)	1.49 (.93)	1.83 (.87)	1.92 (.91)	1.49 (.93)
DD	1.59 (.96)	1.75 (.93)	1.59 (.93)	1.45 (.93)	2.05 (.83)	1.39 (.91)

Note—In this table DA = direct statement, INF = inference, and IDENT = identification.

Final recall. The final recall of predicates is given in Table 5. The ratio of confusions to correct recall in the before condition was .61. This is more than the ratio from Experiment 1 (.52). So this is further evidence that subjects were treating the individuals in the before condition in a more unified manner because of the use of pictures. The ratio in the after condition was .34, which is identical to Experiment 1.

Discussion

The general conclusion from this experiment is that the use of portraits does promote the formation of unitary impressions in the before condition. However, there was an inference effect, if reduced, in the before condition. This indicates that sometimes subjects are not forming unitary impressions. The other fact to note again about this experiment is the remarkable fidelity with which the present study has replicated the asymmetries in the after condition found by Anderson and Hastie (1974). This supports the notion that subjects handle their predicament in the after condition by abandoning the node set up for the proper name.

EXPERIMENT 3

Experiment 3 had three purposes. The first purpose was to see if the before inference effect could be further reduced. The second was to put to test the explanation of the asymmetries in the after condition. The third was to further trace out the time course of the diminution of the inference effects. The first two purposes were served by changing the labels used. So, rather than using a proper name and a profession, one label was an adjective-plus-nationality (e.g., "tall Russian") and the other an adjective-plus-profession (e.g., "smart lawyer"). Except for this, the procedures and materials were identical to those in Experiment 2.

In fact, subjects saw the exact study sentences and test sentences in this experiment as in Experiment 2, with the exception that an adjective-plus-nationality always replaced the proper name and an adjective-plus-profession always replaced the profession. It is hard to imagine that there should be any difference in ease of learning facts true about a profession vs those about a nationality. Therefore, if the asymmetries in the after condition of previous experiments were due to poor learning of proper name information, then there should be no asymmetries in the after condition of this experiment.

It was also thought that this move would eliminate the inference effect in the before condition. In the first identification phase, the subject learned to recall the two adjectives, the professions, and the nationality to a portrait. In this phase, the four words were randomized and the subject had no way of knowing which adjective would go with which noun. Thus, in later phases of the experiment, the only way he could integrate the adjective and the noun was to refer to the single individual he had set up for the first identification.

There was one other change in experimental methodology. Four passes were run through the second RT to further trace changes in the inference effects. The second two passes were an exact repeat of the first two. Fifteen subjects from the University of Michigan undergraduate population served in the experiment.

Results

In all statistical tests, no differences between professions and nationalities proved to be significant. Therefore, the results were collapsed over this variable. Figure 8 shows the data for the direct statement trues and the falses as a function of point in the experiment. The standard error of "true" judgment times was 76 msec ($df = 525$) and 51 msec ($df = 1,185$) for "false" judgment times. The "true" data were classified according to whether the statement involved a before or after label. The "false" data were classified according to whether the predicate was true of a before or after label. Classification of "false" data by predicate brings out larger differences than classification by label.

During the first RT, subjects were faster with after material than with before material. They also showed much lower accuracy to before trues than to after

Table 5
Recall of Predicates in Experiment 2

	Correct	Confusion	Incorrect
Before			
PN	45%	29%	26%
DD	50%	29%	21%
After			
PN	48%	27%	25%
DD	68%	12%	20%




Figure 8. Reaction time and accuracy to direct statement trues and to falses throughout the course of Experiment 3.

trues. The poorer performance for before materials can be attributed to the fact that the subject learned four predicates about each individual in the before condition, but only two predicates in the after condition. The additional predicates in the before condition caused greater interference.

The "false" reaction times increased from first RT to second RT, Pass 1, but the trues did not change much. This significant interaction [$t(1710) = 3.71$, $p < .001$] replicates the results obtained in past experiments. There was a general decrease in reaction time over the four passes in the second RT. The "false" times remained longer than the "true" times throughout the second RT. However, the differences between the before and after conditions disappeared. Accuracy improved throughout the course of the second RT. There were no significant differences in accuracy among the four conditions in the second RT.

Figure 9 presents the data on the inference effects during the second RT of this experiment. Reported there are the differences between inferences and direct statements both in accuracy and reaction times. The before condition showed a small but marginally significant effect [an average of 91 msec; $t(525) = 1.69$, $p < .05$, one-tailed] that did not appear to diminish over the course of the experiment. There was also a small but significant effect of inference on accuracy in the before condition [an effect of 2.5%; $t(14) = 2.08$, $p < .05$].

There was a large and significant inference effect in the after condition [an average of 356 msec; $t(525) = 6.62$, $p < .001$] and an effect on accuracy of 7.6% [$t(14) = 5.98$, $p < .001$]. The inference effect decreased over time for both the reaction time and the error rates. The linear trend in the reaction times was highly significant [$t(525) = 3.24$, $p < .001$], as was the trend for error rates [$t(14) = 3.34$, $p < .005$]. This suggests that a subject's memory representation changes from one in which predicates directly studied

of the label are stored closer to the label to a representation in which all predicates are equally available from all labels.

It might seem surprising that there was no change in the small reaction time effect for inferences for the before condition. Also, the after reaction time effect appears to asymptote at a value greater than 0. It does not appear that the inference effect completely goes away at any test interval for either condition. This may be related to the phenomenon of verbatim representation of sentences in long-term memory (e.g., see Anderson, 1974; Keenan, 1975). It seems that, with some small probability, subjects can keep verbatim representations of sentences in memory. The representations are similar to physical images of the sentence as studied. Access to such verbatim representations allows subjects the possibility of quicker verifications than if they have to retrieve the sentence's meaning.

Final recall. The final recall data is displayed in Table 6. In that table, NA refers to nationality labels and PR to profession labels. In this experiment, there was a large difference between ratio of confusions to recall in the before condition (47%) and in the after condition (8%). This reinforces the impression from the verification data, that subjects store predicates at a single individual node in the before condition, but at two individual nodes in the after condition. In contrast, the interpretation of Experiment 1 was that predicates were stored at two individual nodes in both the before and after condition. Experiment 2 seems to fall somewhere between Experiments 1 and 3.

Discussion

Figure 10 presents a schematic representation of the information structure that is hypothesized to underlie performance in Experiment 3. There are three major differences between it and the proposals in Figure 5 for Experiment 1. First, in Experiment 3, four labels




Figure 9. Differences between inference and direct statement throughout the second RT of Experiment 3.

Table 6
Recall of Predicates in Experiment 3

	Correct	Confusion	Incorrect
Before			
NA	56%	26%	18%
PR	58%	28%	14%
After			
NA	82%	6%	13%
PR	83%	8%	10%

predicates are being copied from the abandoned node to the other node. As the links involving the abandoned node are weakened because of disuse, the representation for the after condition becomes functionally very similar to the representation for the before condition. Therefore (see Figures 8 and 9), reaction time differences between the two conditions disappear.

CONCLUSIONS

Two major points should be emphasized: First, it is not easy to cause subjects to treat two labels as having the same referent. Use of a portrait drawing (Experiment 2) has a useful role in convincing subjects. Even more powerful, however, is the device used in Experiment 3, of making it cognitively "worth their while" to treat different labels as having identical referents.

The second conclusion concerns how subjects handle their predicament in the after condition. Subjects first encode via a proposition that the two individual nodes turned out to have the same referent. Then they chose to maintain the "stronger" node, the one with more information. They begin a process of copying information from the abandoned node to the saved node. Eventually, the links involving the abandoned node become weakened through disuse and the subject loses access to the abandoned node. The fact that subjects do abandon one node when they discover two nodes have the same referent is evidence for a basic assumption of many semantic network models: the one individual-one node principle.

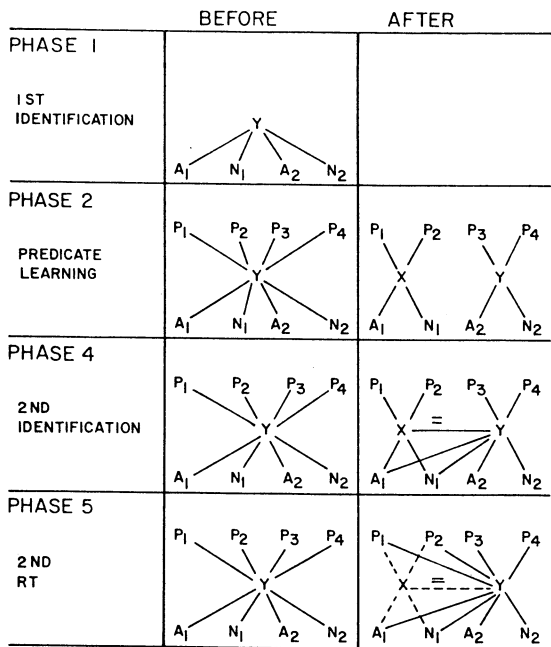


Figure 10. A schematic illustration of the change in information representation over the course of Experiment 3.

rather than two provide access points to the individual nodes. The two adjectives are labeled A₁ and A₂ and the two nouns, N₁ and N₂. Second, a single individual node is used in the before condition. Third, in the after condition, the individual node set up for the nationality is as likely to be abandoned as the one set up for the profession. Therefore, no asymmetries were predicted and none were found.

Figure 10 predicts all the major aspects of Experiment 3. Individual nodes in the before condition in the first RT are attached to four predicates and four identifying labels, while after individuals are attached to two each. Therefore, there is more interference in the before representation and performance is predicted to be worse in the before condition during the first RT. The addition of links to the after representation produces a slow-down relative to the before condition from first RT to second RT, first pass. Differences in RT between the inferences and direct statements decrease throughout the second RT phase because

REFERENCES

ANDERSON, J. R. Verbatim and propositional representation of sentences in immediate and long-term memory. *Journal of Verbal Learning and Verbal Behavior*, 1974, 13, 149-162.

ANDERSON, J. R. *Language, memory, and thought*. Hillsdale, N.J: Lawrence Erlbaum, 1976.

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

ANDERSON, J. R., & HASTIE, R. Individuation and reference in memory: Proper names and definite descriptions. *Cognitive Psychology*, 1974, 6, 495-514.

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.

COLLINS, A. M., & LOFTUS, E. F. A spreading-activation theory of semantic processing. *Psychological Review*, 1975, 82, 407-428.

COLLINS, A. M., & QUILLIAN, M. R. How to make a language user. In E. Tulving & W. Donaldson (Eds.), *Organization and memory*. New York: Academic Press, 1972.

FREGE, G. Über sinn und bedeutung. *Zeitschrift für Philosophie und philosophische Kritik*, 1892, 100 (all). (Translation available in P. Geach & M. Black (Eds.), *Translations from the philosophical writings of Gottlob Frege*. Oxford: Basil Blackwell, 1960. Pp. 56-78.)

KEENAN, J. M. *The role of episodic information in the*

- assessment of semantic memory representations for sentences.* Unpublished PhD thesis, University of Colorado, 1975.
- KINTSCH, W. *The representation of meaning in memory.* Hillsdale, N.J: Erlbaum, 1974.
- LOUX, M. J. The problem of universals. In M. J. Loux (Ed.), *Universals and particulars.* New York: Doubleday, 1970.
- MENDELSON, E. *Introduction to mathematical logic.* New York: Van Nostrand, 1964.
- NORMAN, D. A., RUMELHART, D. E., & THE LNR RESEARCH GROUP. *Explorations in cognition.* San Francisco: W. H. Freeman, 1975.
- RAMSEY, F. P. *Universals: The foundations of mathematics.* New York: Harcourt & Brace, 1931.
- RUSSELL, B. *On the relations of universals and particulars.* Proceedings of the Aristotelian Society, 1911-12, 12, 1-24.
- WINER, B. J. *Statistical principles in experimental design.* New York: McGraw-Hill, 1971.

(Received for publication January 28, 1977;
revision accepted February 3, 1977.)