

SEMANTIC MEMORY RETRIEVAL: SOME DATA AND A MODEL

Elizabeth F. Loftus
University of Washington

Introduction

If I were to ask the average English-speaking adult to name an animal, or a yellow fruit, or a flower beginning with the letter "D", he would ordinarily be able to give me a correct answer in less than a few seconds. Producing a piece of information that one has learned some time ago and knows very well requires very little effort indeed. But, how do we do it? By what process do we reach into the huge store of information in memory and produce a response that is appropriate to a given question. There is no doubt that we are very good at doing this, but at the moment we know relatively little about how we do it.

Although it is too early to provide details of the retrieval mechanism, we can make certain general distinctions among possible types of processes. For example, we can classify processes of memory-access in terms of the extent to which they involve successive as opposed to simultaneous consulting of the memory store. This basic distinction has many important consequences: one of them concerns the effect of size of the array to be searched on searching time. To the extent that the retrieval process involves successive consultation of the memory store, retrieval time should be a function of the number of items in the array to be searched. Conversely, to the extent that simultaneous consultation (or "parallel processing") is involved, retrieval time should be relatively unaffected by the number of items in the array.

Previous Research

In most of the earlier research dealing with the effect of array or category size on retrieval time, subjects have been given a recognition or identification task. Typically they must decide whether a given instance is a member of a particular array, class or category. Most of this work has dealt with relatively small, newly learned categories that are typically defined solely by an enumeration of instances (e.g., Neisser, 1963; Pollack, 1963; Rabbitt, 1959; Sternberg, 1966). For example, Sternberg (1966) gave subjects a group of digits (e.g., 4, 8, 3, 5) and asked them whether another digit (e.g., 3) was a member of that set. This research on "short-term memory" has generally found response time for both positive and negative instances to be a direct function of the size of the category. This result has been interpreted as evidence for the existence of a successive and perhaps even exhaustive scanning process for retrieval from short-term memory.

What happens when larger and better-learned categories are used? For example, suppose we ask a subject to decide whether a "dog" is an "animal"? Landauer and Freedman

(1968) found that subjects took longer to decide that an item was not a member of a large category than to decide that it was not a member of a small category. For example, it took longer to decide that a "desk" was not an "animal" than to decide that it was not a "dog". There was a slight, but nonsignificant, difference in decision times for the positive instances. Collins and Quillian (1969) and Meyer (1970), using somewhat different procedures, also report that retrieval takes longer for larger categories.

Although Landauer and Freedman, Collins and Quillian, and Meyer all find an effect of category size, the effect is very small relative to the size of the categories used. For example, Landauer and Freedman report that the correct identification of negative instances took 53 msec. longer for large than for small categories. Given the relative size of the arrays to be searched, in order for such small differences to be caused by the additional time it takes to scan the larger categories, the successive scanning would have to be enormously rapid, on the order of 1,000 words per sec., which seems improbable. The authors accordingly conclude that the small differences should probably not be interpreted as evidence for the existence of a successive scanning process. Thus, previous research has shown some effect of category size on retrieval from semantic memory, but it has generally not been interpreted in terms of successive processing. On the contrary, the differences are so small that they seem to argue for a mechanism consisting largely if not entirely of parallel processing of some sort.

The Original Experiment

In 1969 I conducted an experiment on this problem in collaboration with Jonathan Freedman at Stanford University. Since subsequent experiments have used a similar methodology, and have in general been designed to test a model of retrieval that we proposed in the publication growing out of that research (Freedman and Loftus, 1971), it is appropriate to describe the original experiment in some detail.

The experiment was designed to provide more information about the retrieval mechanism with particular emphasis on the question of the extent to which successive processing is involved. One major departure from previous work is that rather than study identification time as had been done before, this experiment concerned the speed with which a subject could actually produce a word himself. Instead of giving him a stimulus and asking him to decide whether it was a member of a category, he was given a category and asked to produce a member of it. The identification procedure has been used in the past largely because it is simple, convenient and yet obviously involves the retrieval of information from the memory store. The production procedure is more complicated and less convenient because in a sense it involves a higher level of retrieval—the individual must not only reach into the store and find information relating to a stimulus, he must actually produce an item from that

store. Despite its difficulties, the use of this procedure seemed to be justified because it is an important form of retrieval and provides additional information about the retrieval process.

The basic procedure was to present the subject with a stimulus consisting of a noun category paired with either a letter or an adjective, and ask him to provide a word that belonged in the overlap defined by the pair. For example, subjects who were presented with the pair "fruit-P" might say "peach", "pear", or "plum", among other possibilities. A correct response would be any word beginning with "P" that names a kind of fruit. Responses such as "apple" or "pony" would be incorrect.

As described above, the main focus of the study was on the extent to which the retrieval mechanism involved successive scanning. The effect of size of the array to be searched on retrieval time was an indication of the presence or absence of successive processing. The effect of array size was investigated by selecting many noun categories that ranged in size from very small (e.g., seasons, gems) to extremely large (words, first names).

Another consequence of successive processing concerns the effect of the array defined by the overlap of the noun category with the adjective or letter restrictor, or equivalently, the number of possible correct responses. The size of the overlap varied from a minimum of one ("animal-Z", with "zebra" being the only correct response in most people's memory store) to very large ("first name-J", with dozens of names being possible responses). To the extent that successive scanning plays a role in retrieval, it seems likely that the larger the number of possible responses, the sooner a subject will find and produce one.

The study thus provides two ways of assessing the extent to which successive scanning occurs in the retrieval process: the effect of size of the category on reaction time and the effect of overlap of the category and modifier on reaction time.

Method

Forty subjects were individually told that we were conducting a study of how memory worked, that they would see items consisting of categories and either adjectives or letters, and that they should respond with a word that was an appropriate member of the category. They were given examples, and told to respond as quickly as possible, but to avoid errors.

During the experiment, the subject sat in front of a screen in which was a window covered by half-silvered glass. An index card containing the stimulus was placed in a dark enclosure behind the mirror and was presented by illuminating the enclosure. A microphone was placed in front of the subject and he responded by speaking into it.

A trial consisted of the following. As a card with the item printed in large type was placed in the darkened enclosure behind the half-silvered mirror, the experimenter said "Ready", and pressed a button that illuminated the first member of the stimulus pair. After an interval of at least .5 sec., the second member of the stimulus pair was automatically

illuminated, and simultaneously an electric timer with a DC clutch was started. The subject's verbal response activated a voice key that stopped the clock and terminated the trial. Each subject received a random sequence of 96 stimuli with the category preceding the restrictor on half the trials and following it on the other half.

Results

The two results that bear on the successive scanning issue are as follows.

- 1) There was a nonsignificant correlation of $-.22$ between category size and reaction time. Thus, contrary to expectations from a successive scanning model, the amount of time to retrieve a word from a large category is not greater than for a small category—rather there is little relationship between size and latency and what there is indicates that larger categories take less time than small categories.
- 2) There was a nonsignificant correlation of $.02$ between the number of possible correct responses and reaction time, also contrary to expectations from a successive scanning model.

Although size of category and overlap do not affect reaction time, other variables do.

- 1) Recall that each stimulus consisted of a noun and either a letter or an adjective; and that the noun came either first or second. Reaction time was significantly faster when the noun came first rather than second (1.87 vs. 2.12 sec.). Reaction time was also significantly faster when the stimulus included an adjective rather than a letter (1.84 vs. 2.15 sec.).
- 2) A second variable involves the frequency in English of the possible correct responses. Specifically the frequency in English of the one correct response that had the highest such frequency was related to reaction time. Stimuli that have higher frequency responses tend to have faster reaction times. This finding is a further demonstration of the frequency-reaction time relationship found by Marbe quite long ago (Thumb and Marbe, 1901; cited in Woodworth and Schlosberg, 1954).
- 3) Another important variable involves the likelihood that a particular response will be given when subjects are asked to name words that fit a particular category. We call this the "dominance" of a response within a category. Rather than the frequency in the English language in general, "dominance" refers to the frequency with which a word is given as an example of a category. Information on dominance was obtained from Battig and Montague (1969). It is clear that high dominance produces faster reaction times. Within the "fruit" category, for example, "apple" is more dominant than "lemon"; thus there is a tendency to respond to "fruit-A" more quickly than to fruit-L".

Discussion

The major focus of this study was the question of whether the process of retrieval from semantic memory involved successive scanning to any appreciable extent. To the degree that

successive scanning of items in the memory store occurs, reaction time should be longer when more items have to be scanned. When a subject is asked to produce an item from an array or category, the larger the category, the longer the reaction time should be. Thus, if reaction times are longer for larger categories, it suggests that successive scanning is occurring; if reaction times are not longer for larger categories, it suggests that successive scanning is not occurring. The present study did not find a positive relationship between category size and reaction time. Furthermore, to the extent that successive scanning plays a role in retrieval, reaction time should be shorter when there are a large number of possible correct responses (the more correct responses, the easier it should be to find one). The present study did not find a negative relationship between the number of possible correct responses and reaction time. Thus, all of the evidence indicates that the process of retrieval involves little or no successive scanning of the memory store.

We might mention at this point that the items being retrieved were extremely well-learned. It may well be that with relatively obscure items, or with easy items that for some reason are not found quickly, a subject eventually resorts to successive scanning. However, it appears that with these overly-learned items, the normal, successful process of retrieval does not involve successive scanning of the memory store.

A Retrieval Model

If the retrieval process does not involve successive scanning, of what does it consist? One model that appears to fit much of the data involves a hierarchical organization, by which we mean a system that is divided into a number of interconnected subsystems, each of the latter being hierarchical itself. Specifically, our basic conception is that of a memory organized into a complex hierarchy composed of categories (e.g., animals) with subsets of each (e.g., birds, dogs) and supersets (e.g., living things). The search process need not begin at the top and work down until the appropriate subset is found (as suggested, for example, by Green, et al., 1963 and Lindsay, 1963). Rather, each category heads its own hierarchy that can be entered directly. Thus, to find a bird the category "birds" is entered; to find an animal, the category "animals" is entered; and so on.

In order to explain the lack of effect of category size, additional assumptions are necessary. Within each category it seems likely that a variety of subsets exist. Some are noun categories that are specialized members of the larger set (e.g., birds and dogs are subsets of animals). Some are clusters of words that are highly associated for any of a number of reasons. They may have qualities in common (e.g., all begin with the letter "S" or all have long names or all rhyme). They may be associated with each other in the individual's experience (e.g., all be animals in the Pooh books). These clusters are probably more idiosyncratic than the main, noun categories, but are somewhat consistent across individuals.

And finally, under each category is an undifferentiated enumeration of exemplars. Within all subsets, clusters and the enumeration, the instances will be listed in a more or less constant order, according to their frequency in the language or dominance in the category.

Using this as a general model, we can begin to describe how a person manages the task that he has been given in the above experiment. He must, for example, find a fruit beginning with the letter "P". It seems reasonable to assume that the process of retrieval has at least two major steps: 1) entering an appropriate category—"fruits"; and 2) finding an appropriate member of that category—"peach" (plus, of course, the time required to produce the response verbally). These two steps may each be divided into substeps, but for the moment we have no information on that.

It can be seen that nothing in the model would imply that retrieval would take longer for larger categories. Step 1 is assumed to be equally easy for large and small categories, since the particular category is entered directly. Similarly, but somewhat less obviously, step 2 need not be affected by category size. The crucial problem is finding an instance of the category that fits the restriction imposed. Once the category is entered, the next step (presumably a substep) is to find an appropriate cluster of items. Under fruit, a cluster of "fruits beginning with P" exists, is entered, and the first entry is read out. If instead the task were to find a gem whose color was green, the smaller set "gems" would be entered directly, the subset "green gems" found and the first entry read out. And to find "Noun-Y", a huge set with a huge subset, the same process is repeated. In all cases, it consists of finding a main set and a subset, and does not involve searching through an array of instances.

This model of the retrieval process accounts for the lack of effect of category size, and, it seems to us, is quite plausible. The one implausibility is the notion that so many clusters exist. It seems unlikely that every category contains separate clusters for each letter and adjective. However, there are two arguments in favor of the existence of such clusters. First, when a subject is asked to name several fruits beginning with a particular letter, he has little difficulty in producing many correct responses, very rapidly (if, in fact, many common responses exist). That is, he does not say "peach", pause and then say "pear", pause and then say "plum". The typical pattern we have found in informal testing is a spurt of several responses, perhaps followed by a pause and then several more responses, etc. This suggests that these words are clustered in some sort of functional relationship. The second argument in favor of the existence of clusters is that when multidimensional scaling and clustering procedures are applied to category retrieval data, certain category members that have qualities in common tend to cluster together (Shepard, 1972).

In this model, we have made the assumption that the memory store is organized primarily into noun categories. It almost certainly includes groups of items that have quali-

ties in common (e.g., warm, red, etc.) and may even have groups that have initial letters in common (e.g., all words starting with P). In fact, we can handle the fact that retrieval time is faster when a noun is paired with an adjective rather than a letter by postulating that the main clusters are items that have qualities in common and within these clusters are groupings of items that have initial letters in common. The major, useful organization, however, is thought to be in terms of noun groupings. The significance of this assumption is that the first step of the retrieval process would consist of locating and entering the appropriate noun category. The present experiment provides some data that are consistent with this model and that provide an estimate of the duration of the first step in the process.

Each stimulus pair was presented with the noun either first or second. When the noun comes second, the total retrieval process begins only after its presentation. When the noun comes first, step 1 can be begun before the second half of the pair, any difference due to order can be assumed to be caused by that part of the process that is completed before the second half is shown. Thus, the difference between noun first and noun second is an indication of the existence of step 1 and of its duration.

When nouns are presented first, the mean reaction time is 1.87 sec., with noun second it is 2.12 sec. This difference of .25 is quite stable across somewhat different conditions. With noun-letter pairs, it is .27 sec.; with noun-adjective pairs it is .23 sec. These data support the idea that step 1 consists of entering the noun category, and suggest that the duration of step 1 is approximately .25 sec. for the procedure employed here.

To summarize, we have presented a hierarchical storage model, and have proposed that the retrieval process from semantic memory consists of two distinct steps (entering the appropriate category and then finding an appropriate member of that category). This general model is consistent with the lack of relationship between category size and retrieval time, and is given some support by the difference in reaction times when the noun is presented first and second. That difference of .25 sec. also serves as an indication of the duration of step 1 of the retrieval process.

Experimental Support for the Model

Now let me present some evidence supporting three different aspects of the model.

1) The semantic memory store is organized primarily into noun categories.

The assumption that the major organization is in terms of noun categories has several important implications. One observable consequence concerns the number of words that a subject can produce in a given amount of time under various conditions. If we present a noun category and give a subject, say, one minute to write down any items that belong in the category, he should be able to produce more items than if we present him with an adjective and ask him to write any items which have the inherent quality of that adjective. According to the

model, when a subject is presented with a noun category, he enters that category in his semantic memory and begins to name members of the category that are stored there. When he sees an adjective, however, he must still enter the memory store at some noun category, whereupon he can name members of that category which have the inherent quality of the presented adjective. When he has names as many members as he can, he must shift to another category and begin to name items from it. Thus, when presented with an adjective fewer items should be produced for two reasons: (a) The subject must decide which category to enter first, and this decision takes time, and (b) he must shift from one category to another whenever he needs to, and shifting takes time

An experiment was designed to determine whether the number of responses given to a noun is different from the number of responses given to an adjective. Two hundred students were told that they would see stimuli consisting of either a noun or an adjective. If the stimulus were a noun category, they were to write down as many items as they could that belonged to that category. For example, subjects who were presented with the stimulus "seafood" might say "oyster", "clam", and "shrimp", among other possibilities. If the stimulus were an adjective, the subject was to write down as many items as he could that had the inherent quality of the adjective. For example, if the stimulus were "hard", the subject might say "brick", "rock", "wood", etc. Details about the procedure can be found in Loftus (1972).

The mean number of responses given in one minute to noun stimuli was 12.03 and to adjective stimuli it was 9.15. More detailed analyses of the response protocols indicated that when the stimulus was a noun, almost all the responses were instances of that noun category. However, when the stimulus was an adjective, subjects tended to name a number of items from one category, shift to another category and name a number of items from that category, shift again, and so on. One subject's response protocol to the stimulus "small" illustrates this tendency: "rat, mouse, canary, parakeet, toe, finger, eye, ear". This subject, it can be seen, named two animals, two birds, and four parts of the body. These results are consistent with a model that describes memory as being organized primarily into noun categories. Only after a subject has entered a noun category does he find clusters of items that have qualities in common. In other words, after completing the first step of the retrieval process (category entry), the subject may find a cluster of instances to which a given adjective applies.

The second piece of evidence in support of the assumption that the memory store is organized primarily into noun categories grew out of a conversation with Allan Collins at a 1971 meeting of the Eastern Verbal Investigator's League. We were discussing the finding that reaction time is significantly faster when a noun category comes before the adjective or letter restrictor. This finding is somewhat surprising in view of the fact that English language

habits favor an adjective-noun order. We talk about a "yellow bird" but never a "bird yellow". The question arises as to whether there are special adjectives such that reaction time in a production task would be faster if the adjective preceded the noun. Perhaps if we could find adjectives that were more closely related to particular instances than were the superordinates of those instances, a reversal would occur. For example, if "sour" is more closely related to "lemon" than is "fruit", then perhaps presenting "sour" first and "fruit" second would result in a faster "lemon".

We used the Kent and Rosanoff (1910) norms as one course of adjectives. These norms are the associations of 1000 subjects to 100 familiar English words. In the Kent and Rosanoff norms, the listing of responses to "sour" indicates that "lemon" or "lemons" is the second most common response. The only word given more often is "sweet". In the Battig and Montague (1969) norms, however, the responses to "fruit" indicate that "lemon" is the ninth most common response. Thus, in some sense, "sour" is more closely related to "lemon" than is "fruit". Even in cases such as these, however, reaction time is shorter when the noun is presented first. The retrieval model proposed accounts for this fact by assuming that memory is organized primarily into a complex hierarchy of noun categories and that the first step of the retrieval process is to enter one of these noun categories.

2. Each category in the hierarchical organization can be entered directly.

A portion of a hypothetical memory structure might consist of "living thing" with "animal" and "vegetable" as subsets of it, the superset "bird" and "snake" as subsets of "animal", and "canary" and "robin" as instances of the subset "bird". Suppose for a moment that a category name cannot be located directly, but must be found by beginning at the top of the semantic hierarchy and searching downward through the hierarchy. It follows that the time required to retrieve an instance of any category should reflect the number of supersets through which the subject must move before finding the appropriate category. Thus, we would expect the time taken to retrieve an instance of the category "bird" to be greater than the time to retrieve an animal, because the subject must move through at least one extra superset to locate the category "bird". Note that we have operationally described the hierarchy by assuming that, for a superordinate-subordinate pair of categories such as animal-bird, the superordinate (which includes everything that belongs in the subordinate category) will be higher on the hierarchy. For any superordinate-subordinate pair, then, retrieving an instance of the superordinate category should take less time, according to the view that each category cannot be located directly.

Alternatively, suppose that when a search process begins with a particular category name, that name is directly accessible by some central processor, instead of being accessible only via a search along a hierarchical path. In this case, there is no reason to expect

any difference in the time taken to retrieve an instance of the category "bird" and the time taken to retrieve an animal.

We have designed an experiment to determine whether the time required to name an instance of a category is dependent upon the position of the category in a semantic hierarchy. In designing the experiment, a central problem was to ensure that categories used could be classified accurately according to their relative position in the hierarchy. This was accomplished by using pairs of nested categories in which the superordinate category included by definition everything that belonged in the subordinate category. For example, the category "beverages" contains all alcoholic beverages, "musical instruments" contains all stringed instruments, and "animals" contains all birds. In this experiment the subject was presented with both a superset name (e.g., beverage), and required to give a member of that category, and a subset name (e.g., alcoholic beverage) and required to give a member of it. This allowed a direct comparison of the time taken to produce members of supersets and subsets. For additional details, see Loftus, et al. (1970).

Subjects took an average of 1.60 sec. to produce a member of a superordinate category, and 1.49 sec. to produce a member of a subordinate category. This nonsignificant difference is inconsistent with the view that a search process begins at the top of the hierarchy and follows pathways downward through the network. Such a view would predict longer times for a subset category than for a superset category since the subset is located closer to the bottom of the hierarchy. The present result actually indicates a slight difference in the opposite direction. These data argue in favor of a model in which each category can be located directly.

3. The retrieval process consists of two major steps.

In this section, we discuss two experiments that were originally designed to provide more information as to why some items are responded to quickly while others require a considerably longer response time. For example, subjects take less than 1.00 sec. to name a color, but they take nearly twice as long to name a building (Loftus, et al., 1970). If each category is entered directly, the reason that some categories are less accessible than others cannot be due to the necessity of searching through more of the hierarchy to reach those categories.

One possible answer to the question of why the speed of producing instances of some categories is faster than others is suggested by the effect of word frequency on identification time. It is well established that reaction times to identify high frequency words are faster than for low frequency words (Rubenstein, et al., 1970). This difference in ease of recognition may account for difference in the speed of producing instances of a category.

Clearly, stimulus recognition is a logically necessary step in production of a response from memory. In other words, when a subject is asked to name a member of a particular category, he must "recognize" the stimulus word (in this case, the category name) before he can find and produce a member of the category. Thus, the production of the response "peach" to the stimulus "fruit" depends not only on the existence in memory of a link between "peach" and "fruit", but also upon the semantic identification of "fruit" when it is presented. The recognition or semantic identification of a category name may be equivalent or at least a part of what we have been calling "category entry" or step 1 of the retrieval process.

If the first step in the retrieval process is indeed recognition of the category name and if recognition speed is a function of the frequency in the language of that name, then this step of the process should clearly take longer for categories whose names are low frequency words. Other things being equal, then, the total time to name a member of a category with a high-frequency name should be shorter than the time to name a member of a low-frequency category.

An experiment was conducted in which subjects were presented with a series of category names, and had to respond with the first word that they could think of belonging to the category. We will refer to this as Experiment I; more complete details can be found in Loftus and Freedman (1972). The mean time to produce a member of a category whose name is a high frequency word was 1.64 sec. and to produce a member of a category whose name is a low frequency word was 1.89 sec. The difference was highly significant, indicating clearly that the frequency of the category name is associated with speed of producing an instance of that category. We have suggested that this is due to the greater time required merely to identify the lower frequency name. The possibility remains, however, that the frequency of the name also affects the second part of the retrieval process—that part which occurs after the name has been identified. In other words, once the stimulus has been fully "recognized", it may still take longer to name a member of a category with a low frequency name. A second experiment, which we refer to as Experiment II, was designed to test this possibility.

In the second experiment, the problem was to design a procedure that would eliminate differences due to the first (reading and identification) part of the retrieval process but that would enable us to measure the speed of the rest of the process (finding and producing the correct response). The first requirement was met by giving the subject the category name a full 3 sec. before he had to make any response. Since the first step takes considerably less than 3 sec. presumably this provided ample time for the subject to read and identify any category name, and thus eliminated any differences in response time due to difference in reading speed. It was essential at this point to prevent the subject from immediately begin-

ning the rest of the retrieval process until we in essence told him to start. This was accomplished by requiring a variety of responses to the category name so that the subject did not know ahead of time which response he would have to give and accordingly would not know for which to begin searching. The idea was that under these circumstances the subject would wait until he was asked for a particular response to begin the appropriate retrieval process.

The subject was given a category name and was required to answer several questions about that category. The questions (instructions) were: (a) "first letter", in which the subject responded with the first letter of the category name; (b) "last letter", in which case the subject responded with the last letter of the category name; (c) "length", in which case the subject responded with the number of letters in the category; and (d) "member", in which case the subject responded with the first word that he could think of that named an object that belonged in the category. We took precautions to ensure that the subjects would not always expect to produce a member of any given category. For details the reader is referred to Loftus and Freedman (1972).

The data of interest were the response times of correct responses to the "member" question. The mean time to produce a member of a category whose name is high frequency was 1.48 sec. and to produce a member of a category whose name is low frequency was 1.67 sec., a highly significant difference. Thus even after the subject has had ample time to read and identify the category than a high frequency category.

What do we make of these results? Recall that in the original experiment, subjects were presented with a noun category plus a restricting letter or adjective, and produced an instance of the category which satisfied the restriction imposed. When the category name was presented first, the mean reaction time was 1.87 sec.; with the category name presented second, it was 2.12 sec. The difference of .25 sec. was quite stable across somewhat different conditions, and was taken as evidence that the duration of category entry, or step 1 of the retrieval process, was approximately .25 sec.

The two experiments just discussed also provide an estimate of the duration of step 1. In Experiment 1 the mean reaction time was 1.89 for low frequency categories and 1.64 sec. for high frequency categories. In Experiment II the corresponding times were 1.67 sec. and 1.48 sec., respectively. Since Experiment I included identification time (i.e., included category entry) while Experiment II did not, the differences in reaction time between the two experiments provide some indication of identification time. For low frequency categories the difference is .22 sec. (1.89 - 1.67) and for high frequency categories it is .16 sec. (1.64 - 1.48). As expected, identification or entry of a low frequency category takes longer than for a high frequency category. The mean for all categories is .19 sec. Thus, although the

present procedures are quite different from that used in the original experiment and involve relatively simple (unrestricted) responses, the estimates of the duration of step 1 obtained in the two papers is quite comparable (.25 sec. vs. .19 sec.).

Conclusion

Our basic conception of human memory is that of a store consisting of a large number of interconnected and cross-referenced associative and category networks. According to the model, memory is organized into a complex hierarchy composed of categories (e.g., animals) with subsets of each (e.g., birds, dogs) and supersets (e.g., living things). Furthermore, when a search begins with a particular category name, that category can be entered directly, without first searching through other categories. Thus, to find a bird, the category "birds" is entered; to find an animal, the category "animals" is entered, and so on. A general model of the retrieval process assumes that the entire process consists of at least two major steps: (a) entering an appropriate category and (b) finding and producing an appropriate member of that category.

We have presented the results of several experiments to support various aspects of the model. For example, support for the view that the retrieval process consists of at least two major steps comes from the reduction in total reaction time obtained when a subject has already read and "identified" the category name. This reduction ranges from .19 to .25, depending on the task, and it serves as an indication of the duration of the first step of the process.

Further development and testing of this model is badly needed. At this point the model is more like a loose outline than a specific theory. Much more research will be needed before we can actually translate this outline into a concrete model. Even more research will be necessary before we can build a realistic model that truly describes the characteristics of human memory. One fruitful direction would involve a refined analysis of data from individual subjects rather than analysis of data averaged across subjects, because the hierarchical organization surely varies according to individual usage and experience from subject to subject. We plan to turn to such individual analyses in future work.

References

- Battig, W. F. and W. E. Montague (1969), "Category Norms for Verbal Items in 56 Categories: A Replication and Extension of the Connecticut Category Norms", Journal of Experimental Psychology, 80, 3.
- Collins, A. M. and M. R. Quillian (1969), "Retrieval Time from Semantic Memory", Journal of Verbal Learning and Verbal Behavior, 8, 240-247.
- Freedman, J. L. and E. F. Loftus (1971), "Retrieval of Words from Long-Term Memory", Journal of Verbal Learning and Verbal Behavior, 10, 107-115.

- Green, B.R., Jr., A.K. Wolf, C. Chomsky, and K. Laughery (1963), "Baseball: An Automatic Question Answerer", in E.A. Feigenbaum and J. Feldman (eds.), Computers and Thought, New York: McGraw-Hill Company.
- Kent, G.H. and A.J. Rosanoff (1910), "A Study of Association in Insanity", American Journal of Insanity, 67, 37-96.
- Landauer, T.K. and J.L. Freedman (1968), "Information Retrieval from Long-Term Memory: Category Size and Recognition Time", Journal of Verbal Learning and Verbal Behavior, 7, 291-295.
- Lindsay, R.K. (1963), "Inferential Memory as the Basis of Machines which Understand Natural Language", in E.A. Feigenbaum and J. Feldman (eds.), Computers and Thought, New York: McGraw-Hill Company.
- Loftus, E.F. (1972), "Nouns, Adjectives and Semantic Memory", Journal of Experimental Psychology, 1972, 96, 213-215.
- Loftus, E.F. and J.L. Freedman (1972), "Effect of Category-Name Frequency on the Speed of Naming an Instance of the Category", Journal of Verbal Learning and Verbal Behavior, 11, 343-347.
- Loftus, E.F., J.L. Freedman, and G.R. Loftus (1970), "Retrieval of Words from Subordinate and Superordinate Categories in Semantic Hierarchies", Psychonomic Science, 21, 235-236.
- Meyer, D.E. (1970), "On the Representation and Retrieval of Stored Semantic Information", Cognitive Psychology, 21, 242-300.
- Neisser, U. (1963), "Decision-Time Without Reaction Time", American Journal of Psychology, 76, 376-385.
- Pollack, I. (1963), "Speed of Classification of Words into Superordinate Categories", Journal of Verbal Learning and Verbal Behavior, 2, 159-165.
- Rabbitt, P.M.A. (1959), "Effects of Independent Variations in Stimulus and Response Probability", Nature, 183, 1212.
- Rubenstein, H., L. Garfield, and J.A. Millikan (1970), "Homographic Entries in the Internal Lexicon", Journal of Verbal Learning and Verbal Behavior, 9, 487-494.
- Shepard, R. (1972), "Some Illustrative Applications", Handout prepared for Multidimensional Scaling Workshop, University of Pennsylvania, June.
- Sternberg, S. (1966), "High-Speed Scanning in Human Memory!" Science, 153, 652-654.
- Woodworth, R.S. and N. Schlosberg (1954), Experimental Psychology, New York: Holt, Rinehart, and Winston.

Acknowledgment

The preparation of this paper was supported by the National Institutes of Health under Grant No. MH-20280.