

DOCUMENT RESUME

ED 162 262

CS 004 458

AUTHOR Williams, Michael D.
TITLE The Process of Retrieval from Very Long Term Memory. Technical Report No. 7801.
INSTITUTION California Univ., San Diego. Center for Human Information Processing.
SPONS AGENCY Advanced Research Projects Agency (DCD), Washington, D.C.; National Inst. of Mental Health (DHEW), Rockville, Md.; Office of Naval Research, Arlington, Va. Personnel and Training Research Programs Office.
PUB DATE Sep 78
CONTRACT N00014-76-C-0628
GRANT MH-15828
NOTE 44p.

EDRS PRICE MF-\$0.83 HC-\$2.06 Plus Postage.
DESCRIPTORS *Cognitive Processes; *Information Retrieval; *Memory; Problem Solving; *Recall (Psychological); *Thought Processes; Visualization

ABSTRACT

In an investigation into the process of retrieval from very long term memory, four subjects who had been cut of high school from 4 to 19 years were asked to think aloud while attempting to recall the names of their high school classmates. The retrievals were found to be characterized by overshoot, systematic hypothesizing, fabrications, the establishment of search contexts, self-corrections, and the use of a number of basic search strategies. These phenomena, as well as an array of traditional memory phenomena, can be understood on the basis of an interpretation of retrieval as a reconstructive problem solving process. Information about the target item is used to construct a description of some aspect of the item, which is used to recover a new fragment of information about the item; from this information a new description can be formed to retrieve still more information, and so on until the piece of information sought is recovered. Three subprocesses may be identified: finding a context for the search, searching, and verifying that what has been recovered is about the target item.
 (Author/GW)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

ED 162262

U.S. DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRO-
DUCED EXACTLY AS RECEIVED FROM
THE PERSON OR ORGANIZATION ORIGIN-
ATING IT. POINTS OF VIEW OR OPINIONS
STATED DO NOT NECESSARILY REPRE-
SENT OFFICIAL NATIONAL INSTITUTE OF
EDUCATION POSITION OR POLICY

The Process of Retrieval
from Very Long Term Memory

Michael D. Williams

Center for Human Information Processing
University of California, San Diego
La Jolla, California 92093

Report No. 7801
September 1978

The views and conclusions contained in this document are those of the author and do not necessarily reflect the policy of any agency of the United States Government. The research was supported by the Advanced Research Projects Agency and the Office of Naval Research and was monitored by ONR under Contract N00014-76-C-0628. The report is approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Michael D. Williams is now at the Navy Personnel Research and Development Center, Code 304, San Diego, California 92152.

8576005

Acknowledgement

Partial research support was provided by the Office of Naval Research and the Advanced Research Projects Agency, monitored by ONR under contract N00014-76-C-0628. Support was also provided by National Institutes of Mental Health grant MH-15828 to the Center for Human Information Processing. I thank Don Norman, Dave Rumelhart, and Al Stevens for their critiques and helpful suggestions during the course of this work.

The Process of Retrieval
from Very Long Term Memory

Michael D. Williams

In order to capture some insight into the process of retrieval, I have studied a naturalistic recall situation: the recall of names of high school classmates of many years ago. Such a recall process can be a complex and convoluted task taking place over the course of weeks. Incidental details come to mind almost before the recovery of the first name. The high school buildings, the classes attended, the activities engaged in all seem to be a part of the process of retrieval. The following is a sample protocol from a subject asked to think aloud while attempting to recall the names of his classmates from high school:

The first thing that comes to mind is . . . I mean it's almost like images of different snapshots of my high school. You know, I can think of my general science class, and waiting in the lunch line, and halls. Umm. Sort of, Jeff Thompson! He was a friend of mine. Sort of pops into mind and I think umm, we used to stand in lunch line together, and he was in my general science class. That's where I first met him, my freshman year. There was, umm, let's see I'm trying to think of people I interacted a lot with. And some of them . . . are sort of people I've known after high school. Like Bill Newell. I . . . I lived with him for a while, in Portland. Umm, after school. So he sort of comes to mind immediately too. Umm, let's see. I mean I guess it's almost easier for me to think of my home town, and think of people . . . that . . . I've still run into, on occasion, when I go back there. And then sort of check to see if they meet the requirements. Like were they in high school with me. And I can think of people like Buddy Collender, and John Tremble, who still both live in my . . . Ah . . . home town. Umm, . . . I guess it's. It also seems that I want to think of, sort of, it's clear that I have to think of some other situations. It's like I want to think of, sort of prototypical situations and then sort of examine the people that were involved in those. And things like P.E. class, where there was . . . Ah . . . Gary Booth. Umm, and Karl Brist, were sort of, we always ended up in the same P.E. classes, for some reasons. Umm, . . . I can think of things like dances. And I guess then I usually think of . . . of girls (chuckle). Like Cindy Shup, Judy Foss, an Ah . . . Sharon Ellis. I went to grade school with her. Umm, . . . I mean it's sort of like I have a picture of the of the high school dance. You

* All names, except those of public figures and a few cases where the phonetics of the name are at issue, have been changed in the protocols presented in this paper.

know, and these are the women I knew then. Umm, . . . There's a girl I can think of, I have a very good picture of her but I can't remember her name. I can even remember she was, she was from Gales Creek. Ah . . . And Ah . . . I mean I can see what she looks like. I just can't remember her name. She lived near John Randolph, who was another friend of mine.

What is going on in a search of this nature? Do the pieces of incidental information about the school, or where someone lived, or the activities the subject engaged in serve some purpose? What guides a search of this nature? What is the relationship between searches of this nature and the vast array of laboratory phenomena which psychologists have studied for so long?

Methods

The basic approach taken in these studies is observational as opposed to experimental. Rather than manipulating some set of variables to prove one hypothesis over another, I have chosen to perform a detailed set of observations. Subjects are asked to think aloud as they try to recall the names of their classmates in high school. The data are the protocols that the subject generates, including, of course, the names themselves and the times between names.

I had available school yearbooks for the subjects, so that most recalls could be verified. In addition, at the conclusion of the experiment (usually after about 10 hours of recall spread over several weeks), I conducted extensive debriefings of the subjects, attempting to determine the origins of errors.

In the discussion section I present an information processing analysis of remembering as a problem solving process. The analysis provides a theoretical background with which to interpret the observations from the subjects' verbal protocols, as well as an array of traditional memory phenomena.

Many portions of the analysis presented are not new to the memory literature. What is new, I believe, is the wide range of the phenomena being dealt with, and the character of the interactions of various information processing constraints which provide an account for some of the flexibility and apparent limitations of human memory. The final product is a characterization -- a first pass theory -- of the retrieval process which provides an integrated framework for interpretation a great deal of what we already know about memory and exploring new characteristics.

Procedure

Four subjects were studied individually. They were instructed, to think aloud while attempting to recall the first and last names of their classmates from high school. The yearbooks from each subject's high school were available to me so that I could get independent verification of the subjects' recollections. Subjects were run from 4 to 10 hours in approximately one hour sessions. All sessions were tape recorded.

Subjects were first asked a warmup question in order to insure that they understood the protocol instructions. The question was "Name the state capitols that begin with the letter 'B'." During their protocols, subjects were prompted whenever they had prolonged silences with comments from the observer such as: "What are you thinking now?," "What is passing through your mind?," etc. Subjects S1 and S3 were run in sessions of approximately 1 hour, 1 session per day, 5 days per week, for two weeks. Subjects S2 and S4 were run on sessions of from 1 to 2 hours, at more erratic intervals due to their personal schedules.

Subjects were directed to avoid thinking about their high school classmates between sessions, but this was not completely possible. Names were occasionally recalled inadvertently between sessions. Before each session subjects were asked to report all the names they had recalled between sessions and to relate the circumstances of the recalls.

Subjects

The four subjects (S1 through S4) were from 4 to 19 years out of high school. All attended high school in the San Diego area. The subjects were paid \$2 per hour for their participation in this experiment. Individually, their backgrounds are as follows:

Subject S1. S1 is a female, 5 years out of high school. She has lived in San Diego in the same house for her entire life. She attended only one high school for 2 1/2 years, graduating 6 months ahead of her class. Most of the subject's classmates from junior high also went to high school with her. Her graduating class had 609 members listed in the senior yearbook.

Subject S2. S2 is a female, 19 years out of high school. She has lived in the San Diego area except for 4 years of

1. There was one exception to the general procedure. Subject S4 was allowed to tape record her thinking aloud without the presence of the observer. Long silences did occasionally occur with S4.

4

college. She presently has a home in the neighborhood of her high school. The subject attended only one high school, and most of her classmates from junior high attended high school with her. Her graduating class had 318 members listed in the senior yearbook. S2 graduated in the same class as S4.

Subject S3. S3 is a female, 4 years out of high school. She has lived in several different regions of the San Diego area as well as having spent her early years out of the state of California. She attended her first year in one high school, transferring to a second high school for her last two years. Most of her classmates from junior high attended a different high school than the subject. Her graduating class has 750 members listed in her senior yearbook.

Subject S4. S4 is a female, 19 years out of high school. She lived in the San Diego area until after high school. She left San Diego for around 9 years, and returned about 10 years ago. She was a member of S2's graduating class in high school, is a personal friend of S2's, and has maintained contact with S2 throughout the years. S4 attended only one high school, and most of her junior high school classmates attended the same high school. Her graduating class had 318 members listed in the senior yearbook.

Scoring a Protocol

The scoring of protocols is a complicated procedure. Scoring is not a simple matter of listing all the names recalled and verifying them against the subject's yearbook. Subjects expressed varying degrees of certainty about their recollections. They changed their minds (in some instances several times). On occasions they had difficulty verifying the correctness of a name. They mentioned names in passing which they did not intend to be candidates for names of high school classmates. Often subjects remembered nicknames of classmates which are difficult, and sometimes impossible, for the observer to verify. Sometimes subjects confused the name of one classmate with that of another.

There are two basic categories of names mentioned, those mentioned for possible "candidates" as classmates and those mentioned "in-passing." An example of how a name might be mentioned in-passing would be:

um well, my next impulse is to go to the art class where we all -- my teacher's name was Bill Dane, uh. This was in a class where I related to people like human beings instead of other desks, where . . .

In-passing names are operationally defined as those names which the subject identifies as not being a candidate classmate either before the name is mentioned, or within the sentence in which the name is first mentioned. On occasion, identification as a non-candidate name is indicated by the context or tone in which the name is mentioned. All names which are not in-passing names are considered candidate names.

Candidate names can be correct or incorrect as verified by the subject's senior yearbook. I refer to incorrect names as fabrications. Both correct names and fabrications can have three levels of evaluation, "yes it is a correct name" (Y), "no it is not a correct name" (N), or "don't know if it is a correct name or not" (DK). Note that initially the subject must evaluate the name as a Y or DK (or else the name would be categorized as a name in-passing, or not mentioned by the subject at all). However, some subjects spend a good deal of time evaluating names and frequently change their minds.

Sometimes subjects recalled names that they had recalled before. Frequently, this is simply a part of the search process, the recalling of the old names sets the context for the search for names not yet recovered. On occasion, however, subjects recalled a name unaware that the name had already been recalled. Indeed, the judgement of whether a name has already been recalled is itself a recollection that can be correct or incorrect with the three levels of evaluation mentioned above. If a subject already mentioned a name, I refer to it as "old," independent of the evaluation that the subject has assigned.

Another scoring problem is the verification of nicknames. For the purposes of scoring nicknames, I have taken the following policy. If there exists a verified classmate with a proper first or middle name which is commonly given the nickname the subject has specified (e.g., William for Bill, Patricia for Patty), then the nickname is scored as correct. If the subject later mentions the correct proper name in conjunction with the nickname (e.g. "Oh,...Bobo Richards is probably listed as Barbara Richards in my yearbook. That was her 'real' name."), then the nickname is scored as correct. If the nickname can be supported by other pictures or signatures in the yearbook (e.g., Kathleen Johnson is also listed as Nancy Johnson in a separate picture), then the nickname is scored as correct. All other cases are scored as fabrications.

Fabrications

The incorrect names, or fabrications, that subjects mention can be of a variety of types. The categories of fabrications are discussed in detail in the section on recurrent phenomena. Fabrications are difficult to score, in part, because subjects frequently introduced candidate names which were quickly rejected by the subjects themselves. I have set as a scoring criterion

that a subject must entertain an incorrect name for at least 10 seconds as possibly correct (Y or DK) in order for the name to be scored as a fabrication.

Results

Basic Results

The basic results of the scoring discussed above are plotted as cumulative graphs against time in figures 1, 2, 3, and 4 for subjects 1 through 4, respectively. Correct names are defined as all candidate names which were verified by the appropriate year-book. Fabrications on these graphs are defined as all incorrect candidate names which the subject did not reject (i.e. evaluate in the N category) within 10 seconds of the introduction of the name. Dotted lines mark the beginning and end of sessions. The jumps in portions of the curves at session boundaries are due to names subjects reported having recalled between sessions.

Even a brief glance at these graphs tells us two things. One, subjects can go on recalling new names for extended periods of time; even after 10 hours (spread over two weeks time) subjects can recall new names. Two, the number of retrievable names can be extremely large.

I examined the character of the memory search protocols at two levels. The first is a description of the recurrent phenomena that occur during the search processes. I believe that these recurrent phenomena will be found in almost all memory tasks, and indeed, I argue later that there is some evidence for such processes in most tasks that have been formally studied by psychologists. The second level is an examination of the specific search strategies used by subjects in my task. While some of the characteristics of the search strategies used have interesting implications for the retrieval process, the specific strategies used are undoubtedly highly determined by the specific task being performed.

The Recurrent Phenomena

Though subjects differed in the individual details of their recalls, a number of search phenomena recurred throughout their protocols. Subjects searched for names from contexts of specific locations or activities. They built up large bodies of facts about specific individuals or contexts before finding a name or beginning a search within a context. They systematically tested sequences of hypotheses seeking a match that would lead to more information. They overshot their original goal and continued to recover information about individuals even after they had recalled the names. They made extensive use of partial recalls. They made a broad variety of errors and frequently corrected themselves. Though most claimed after the first few minutes of the task that they could not recall any more names, with greater

S1




Figure 1. Correct names and fabrications recalled by S1.

S2




Figure 2. Correct names and fabrications recalled by S2.

S3




Figure 3. Correct names and fabrications recalled by S3.

S4




Figure 4. Correct names and fabrications recalled by S4.

effort all recalled new names for many hours. These phenomena appear over and over again, hence the name recurrent phenomena. Let us now examine the recurrent phenomena; partial recalls, incidental recalls and search, contexts, extended retrieval, systematic hypothesizing, inferential recalls, overshoot, fabrications, self correction, distractions, and reminiscence.

Partial Recalls. Subjects frequently recalled fragments of the information they were seeking. For example, a subject might recall that a particular person's name began with a "D" and that it was a two syllable name. Indeed, a great deal of the process of recalling a classmate's name can be typified as a reconstruction from a variety of such bits and pieces of information. This phenomenon has been alluded to in a variety of experiments as partial recall (see Bernbach, 1970, and Woodworth, 1938, pg.37).

Incidental Recalls and Search. Perhaps the most obvious two observations were the extensive amount of incidental information that was recovered and the search that such incidental recalls produced. The nature of the incidental recalls and the characteristic search strategies that developed are discussed in the section on search strategies.

Contexts. A salient feature throughout the protocols was the extensive use of locations or activities as contexts within which to search for acquaintances. Names were recovered as subjects thought about who "was in my 10th grade art class," or "about the time [the members of a rock band] were playing over at my house." The following is typical of the use of contexts (the parts identified as contexts are underlined):

S1: I was trying to think of Kurt's last name, but I just can't think of it. Umm, okay, let me see if there's any other neighborhoods that I haven't gotten to, that I can remember where people my age lived. Umm . . . hmmm. There is no one that lived way up on the end . . . And now I'm trying to think of the Sunset Cliffs down on Cal Western because a lot of people always used to go there and go tide pool picking and just run around and go surfing. I'm trying to think of all the people that perhaps went surfing or even tide pool picking, that were in my grade. Um . . . if I could see them lined up against -- there's this one cliff down at Newport Beach they always use to line up with there boards and sit down and look at the waves, and then I go down the row and see if there's anybody that I haven't already named. There's John Culver'son. I already named him. And Rod Hackbart, and they use to go surfing, and um there are a lot of older people too. Um, Jim Nelson, I already named them, all those guys used to go surfing. Um, he was older, he was older, and older. He was younger. A lot of these guys were older. Let me see, him and him . . . Okay I was just

going down the list and I don't see anybody that I haven't already seen and there was this one girl who always use to be down there, but she was younger. I already named the people she hangs around with. Um, is there anybody else that I know that use to

Extended Retrieval. Subjects frequently encountered a situation where a person was identified but not completely named, e.g. "I remember Bill somebody . . ." When this happened the subjects recalled additional details about the person. I call this extended retrieval. For example,

S: Bob Peterson, uhh Jerry Paulson. Uhh this is a rock band.

E: What are you thinking about now?

S: I'm thinking about a time when they were playing over at my house. umm There's uh . . .

E: What? What is --

S: I'm try. I'm -- Okay I was imagining the whole room and I was imagining the instruments set up and I'm trying to remember the name of this guy -- who used to do art, and he was also in my 10th grade art class which would also bring a whole lot of people to -- first on that -- what's his name now? Let's see -- [whistle] I'm trying to -- remember his name. At his house was the first time I heard a Jefferson Airplane album. Umm plays the bass guitar, really a strungout looking dude, uhh wow --

Frequently such an attempt at extended retrieval would result in the successful recovery of a name.

Systematic Hypothesizing. As the difficulty of recall increased subjects begin to demonstrate instances of systematic hypothesizing. They generated a sequence of hypotheses, attempting to use the hypothesized information to find some more information. For example, "His first name began with a 'D'. Don, Dave, Douglas,"

Inferential Recalls. In many cases subjects could not explicitly recall a particular piece of information. In such cases, they often attempted to infer the information. I refer to this as inferential recall. The following is an example:

Somebody, . . . I remember this girl who used to play the oboe, and it was junior year, she was our age. Or was she older? I can't remember if she was older or not.

I'm sure she ... now I'm going back to the year book, I'm trying to check to see if her face was in the year book that was ahead of us, the class of 1971. It seems to me as though she might have been next to the one that the year I graduated from. That I was in. I can't. The different things that we wore, starting when we were, once again this goes back to the year-book. When we graduated, these girls wore the scalloped tops. You know, black. And all the years before that they wore a V-neck sort of top (when we went to the studios and recorded that for graduation pictures). And I'm trying to remember what she wore. Was it scalloped or V? And that would determine what yearbook she was in. I'm pretty sure it was scalloped. Was it? She had a very broad smile. Blonde, short blonde hair, white face, what was she wearing. I think she was wearing a scallop. If she was wearing a scallop that means she graduated with us. And I can't remember her name. And I should remember her name too because she was the only oboe player in the whole orchestra. And I can't remember what her name was. I remember her face and I ... now that I look, and try to imagine really hard I think she had the scalloped, so she must have graduated with us. Now I don't really remember her graduation per se. She may have left the school before she graduated from Clairmont and graduated from some other high school. I can't remember. I think she left early. But how could she be in the yearbook if she left early. [Laugh] That was during junior year that she was in the band. I don't know. Now I can't decide where to place her. I just remember her distinctly playing in the orchestra and at least being our age or a year older. And she did leave the school before she graduated. I can't remember her name. Okay. Let's get off of that. Let's go to someone else. Still in the bandroom. Who else was playing.

The basis for this particular inference was wrong: scalloped drapes were worn by the female graduates of the subject's high school the year before she graduated.

Overshoot. Sometimes when a subject recalled the name of a friend, they continued to recall additional information. For example, "Steve Jones, he was a box boy at the Food Basket [store] on the corner, and was on the football team." I call this phenomenon overshoot. Although on occasion overshoot may be an artifact of conversational constraints in the protocol situation, on some occasions it is clearly related to some sort of verification process. Although a name may be recalled, it must still be determined that the name belongs to the person being thought of, and that the person being thought of attended the correct high school in the correct year. For example, in the following

protocols overshoot is used overtly as an attempt to verify a name. (the overshoots are underlined)

S3: [mumbles] Okay. Oh, oh, wait a minute. Now we're back to a class I haven't really thought of before, my world affairs class. There was this guy who used to sit in back of me. And he took Spanish classes and I know his name. His last name began with an "O." Ah ... The name Orin Elliot sticks and that is the first time I've not associated with someone in that history class. Not really a history class, a political science class/
He sat in back of me. he had dark hair. ahhh, I think
 . . . He was Orin Elliot. That was his name.

Or an example where the attempted verification fails:

Carol Hammer, I already mentioned her. Oh and Turner. Linda, Linda Turner. I don't know about her first name, but I know her last name. A blond girl. Yeh, Turner. I don't think that's her name. I don't think Linda is her name. I think it's just Turner, for sure. She wound up in Redlands somewhere. I remember that. I don't remember her first name. Linda doesn't quite fit. But I can't think of any thing else that would fit that name. Linda Turner. . . That seems to fit. I don't know. I don't know for sure, but that's the closest I can come. The last name for sure I remember.

Indeed, Linda is not the correct first name.

Fabrications. Fabrications are defined as those names which the subject recalled, but which are not in the subject's yearbook. I have chosen the term fabrication to suggest some of the reconstructive characteristics of these errors.

Fabrications fell into one or more of five basic categories:

- F1. Acoustic errors. F2. Wrong class. F3. Misassignment.
 F4. Missing Picture. F5. Indeterminate origin/Other.

Category F1 items are names which were evident acoustic errors. For example, "Lloyd Chappin" turns out to be Lloyd Chaffin (this was confirmed during the subject debriefing). Category F2 items are the name of students who did not graduate with the subject's class but were in the subject's high school in a different class (as verified by the subject's yearbook). Category F3 items are when the subject assigns one person the name of another individual (this event can only be detected in unusual circumstances, for example, when the subjects catch themselves in an error or in follow up experiments when pictures were used). Category F4 items are the names of people who were in the subject's grade in earlier years, but for whom there are no

records in the 12th grade (This category is only relevant for subject S1, yearbooks for the other subjects had lists of all those who did not get their senior pictures taken). The final category, F5, is fabrications of indeterminate origin. These are names which I could not track down in some independent manner. They may be people who were friends of the subject, but who went to different schools, or one of the other categories of fabrications which I could not trace.

Table 1 shows the number of fabrications in each category for each subject.

Frequently fabrications were events of substantial structure. Many were caught and corrected by the subject, some shortly after the error was made, some a number of days later. While some fabrications appeared to be simple recalls, others were accompanied by extended retrievals and/or obvious uncertainty about the recall (as was the case with many correct recalls). Some fabrications were highly suggestive of the underlying processes which led to an error. For example, the name Carol Ludlow was initially recalled but was corrected several days later to be Judy Wardlow. Karen Ludwig is a friend of Judy Wardlow's and is often recalled as "Karen Ludlow . . . No! Ludwig. Karen Ludwig." The fabrication "Carol Ludlow" appears to be a confounding of the names Karen Ludwig and Judy Wardlow.

A few instances that occurred rarely show the difficulty in the reconstruction of a name from partial information. For example, One subject recalled someone named "Bill, . . . or maybe Bob . . . [who had a last name] like Proxmire, but it wasn't Proxmire." The correct name is Bob Billmire. All of the elements of the name are accessible, but the subject put them together in the wrong order. "Bill" is the first syllable of the last name, not the first name.

Self corrections. On occasion subjects changed their minds. A name was recalled only later to be corrected. For example, "Fred Knight, no . . . McKnight! Fred McKnight. That's it." I call this self correction. This occurs spontaneously without any information or hints provided by the observer.

Distractions. Often times when a subject made a partial recall of say, a name, and was searching for a missing fragment, an item which an obvious error would be retrieved and would then interfere with successful recall. This section of protocol from S4 is an example of a distraction.

It's interesting. Like Larry. I can't remember his last name. Every time I try to think of Larry, I say the name Larry, and then I think of Larry Shepard. I don't, I hardly know Larry Shepard, but I know his name well. He's a professor at UCSD. So, umm, that blocks out Larry in high school, and his last name. I think

TABLE 1
Number of Fabrications Per Category

Category	Subjects			
	S1	S2	S3	S4
Acoustic Similarity	35	5	17	25
Wrong Year	8	10	26	12
Misassignment of Name	10	1	2	3
Missing Picture?*	26	0*	0*	0*
Unknown/Other	<u>38</u>	<u>11</u>	<u>25</u>	<u>56</u>
TOTAL**	108	25	69	90

*Yearbooks for S2, S3, and S4 listed all seniors who did not have pictures taken.

**A few fabrications fell into more than one category.

If I keep recalling Larry's face and his name, Larry, maybe I'll be able to get back into that area of my memory with more strength.

Among the interesting characteristics of distractions are:

1. Distractions appear to interfere with the retrieval process. Even though the distraction is identified by the subject as a false recovery, subjects complain of difficulty in attempting to keep the distraction from reoccurring. 2. On many occasions subjects spent time attempting to identify the source of the distraction. One subject stated that by identifying the distraction more completely, she will have "cleared the debris" and that further recalls will be facilitated. 3. Distractions generally match the target item in a very limited domain. That is, The subject may be trying to recall a red headed person named Fred with a last name beginning with "D" whom she knew originally from elementary school. Nevertheless, she continues to recall the name Fred Dottington, whom she knows to be someone she knew in college and who, incidentally, has blond hair.

Reminiscence. There is no clear limit to the amount that subjects can recall. Even subjects recalling classmates for one hour per day, five days per week, for two weeks, are recalling new names correctly during the 10th hour. I refer to this as reminiscence. Frankly, I gave up in my attempt to drive subjects to their limit of names.

A Comment on the Validity of Protocol Observations

Psychologists have shown substantial concern over the validity of observations derived from protocols (Nisbett and Wilson, 1977). However, I will argue that the observations I have derived from subjects' protocols are no less valid than the traditional counts of the number of correct items recalled in verbal learning experiments. Each of the recurrent phenomena are defined in terms of either a specific category of information recalled or an ordered set of categories of information recalled.

2 Because of the very logic of the notion of recall, if a subject claims to have recalled that a particular name began with the letter "P," then he must have recalled the information in order to report it. That the subject recalled that the name began with

2. For example, a partial recall is the retrieval of a fragment of some item sought, an extended retrieval is a partial recall of a name - or a description of a person - followed by a collection of incidental recalls about the person. The one exception to the claim that the recurrent phenomena are made up of configurations of various categories of information recalled is the description of a distraction which implies the frequent inclusion of a processing statement by subjects (i.e. that the distraction is interfering with the retrieval process).

the letter "P" is no more disputable than if a subject claims to have recalled the word "house" from a recall list.

A second common set of criticisms of verbal protocols is that "thinking aloud" interferes with the process of recall and introduces unknown demand characteristics. I agree with both of those criticisms. However, the very same criticisms can be leveled against any psychological experiment. It is simply that we have grown so accustomed to certain data collection techniques that we fail to notice their flaws. For example, writing the words being recalled from a word list must introduce some interference to the recall process. In addition, the demand characteristics of writing a list of words precludes observing phenomena such as partial recalls, incidental recalls, context recalls, and extended retrievals.³ If I were to have subjects write a list of the names of their classmates, the names might come faster (or slower) but I would lose all access to the intermediate steps in the search process.

The Strategies

A second set of observations I have derived from subject protocol are what I call the search strategies. They have been separated from the recurrent phenomena because I believe that they are dominated by the demand characteristics of the particular recall task I have given subjects. I believe that the recurrent phenomena will show up in any retrieval task where the data collection techniques permit their observation. Certain generalizations of the search strategies (such as one suggested by Keniston and Flavell, (1) may have importance beyond the scope of this task.

A number of search strategies were identified. These were either used by more than one subject or were used by the same subject over an extended period of time. The manner in which subjects employed different strategies had a number of characteristics. Sometimes the same strategy was used on a number of different occasions by the same subject. Some strategies were used for hours by a subject. Frequently a subject would shift from one strategy to another when the second seemed to hold out hope for improved success, and then shift back a few moments later when the second strategy ceased to be productive. I have named the strategies the subjects used General Association, Activities, Locations, Name Generation, and Pictures.

General Association. Every subject used the strategy of general association at one time or another. Indeed, it appeared that most subjects began their recalls with this strategy.

3. Erdelyi and Kleinbard (1978) have conducted an experiment in a modification of a traditional word list paradigm which demonstrates the effects of what I have called reminiscence.

Briefly, general association is described as the strategy of starting with a known person and searching for people who are directly related in some obvious manner (e.g., a friend, cousin, brother or sister). Occasionally a particular person will suggest a group of people (e.g., a social clique). This strategy is generally the first one employed and goes so rapidly that detailed protocols of the process are difficult to obtain. In most cases the use of this strategy is inferred when the subjects report the relationships of a group of names in a post hoc manner. For example:

S: Right. The guy Mel Hagershon that I had a hard time thinking of his last name. I remember his friend was named Mike and I couldn't remember his last name. Fielding, or something like that sticks into my mind. Maybe because I know Mike Fields that isn't related to school but that name also sticks into my mind for him although it may not be right at all. ⁴ Ah, I thought there was something else I thought of. Let me see. Is there anything else. Jonny Faylan. He was also friends with that group, with Mel Hagershon and Mike and I knew him since I was young.

Activities. A second related approach subjects commonly used was the activities strategy. The strategy is initiated by the identification of an activity in which some group of people engaged. The subjects then proceed to attempt to name all the members of that group. Examples of likely groups are, History class, the Baseball team, the Cheerleaders, the Band, the International club, the "people I played tennis with." The following is an example of a subject using the activities strategy:

I'd like to go back to the bandroom again. In the bandroom, what's in an orchestra and who would play with what? Okay. All the violin people I can think of, the gello people, the bass, Kathy. There was a girl I was trying to remember her last name. There are two Kathys who played the French horn. I already mentioned one. Kathy Dolmore. There was another Kathy and I can't remember her name. Kathy. . . It seems like her last name started with a K or K sound. Kathy, oh I can't remember. Oh, Bartholomew sticks out. Not from the bandroom but I can't get the face that goes with it. Wait a minute, maybe that was a Robert Bartholomew. Was that Bob Bartholomew the one I was trying to think of yesterday? Uhmm. He played the horn, trumpet. I'm not sure. Let's see, who else was in the band? Uhmm. I'm still on the bandroom.

4. Note the possible distraction.

Exhausted.

This strategy was used throughout the protocols of all the subjects. Though after several sessions it became less and less frequent as the subjects ran out of new activities to work on.

Locations. The locations strategy is the case where a subject systematically searches a mental map where target items are likely to be recalled. An example of the use of this strategy is the following:

S1: That's a new name. she lived on Alvian Street. Ah, let me see, on the other side of there, there is Bob. He was a year older and that other girl was a year older. OK. So there is Margaret Mott and if you come up and around, there is no one right there. Jim Gott lived down there but I've already named him. I already named Judy Nicholson and let me see if there is anybody on that cross street. There was a girl that was younger and she had a brother that was a year older than me. Then there was a girl a year younger and there is Gay Masterson, I already named her, and Barbie Tollen. They live on the same street and they have no one else on that street that went to Point Loma with us. If I keep going down Silvergate there is someone who lived on the corner house. It had a purple door.....

This particular strategy was used for over 2 hours by S1.

Name Generation. The name generation strategy is where the subjects invent some scheme to generate common first or last names, and then test the names to see if they match the names of high school classmates. One common way of doing this was to proceed through the alphabet, generating common male or female first names. For example:

Are there any other Bettys that I knew? I think she's the only Betty I ever knew. Betty, other girls names with "B"s that are sort of familiar names. Barbara and, I named Barbara Shafer already, and there was Barbie Tollen and I named her. Barbara, Barbie, no those are the only two Barbaras I know. Um, another "A" name for a girl might be, um, there is Ann. I don't know any Alices. Um, no Alissons. Ann. What other "A" name. Let's see, Ann and Alice. No, I don't know anything else like that. Okay, "C"s, or "B"s, are there anymore "B"s. Let's see, there was Barbara and Betty and, um, "B"s. No. Let's see, "C"s. Cathy and I named Kathy Jackson, although her name starts with a "K." Um, Cathy, . . .

This generation of names from the alphabet also frequently was used to produce last names. Another means that subjects used was to look around the room naming objects and generating names from the objects named. For example:

I was just looking at the wheel of, of that chair, and I was thinking of Wheeler . . . Linda Wheeler? That name, I don't know if that was in my . . . Linda Wheeler, that name -- Now there's a name that doesn't have a face that goes with it. I'm not even sure that . . . No! That's not in high school. That's here. That's here in college. Scratch that. That was Lynn Wheeler. She was a roommate of one of the friends I have here on campus. So that's out, she wasn't even anywhere near my school at all. 5

While this later form of name generation occurred in short bursts through out all sessions, the systematic searches using the alphabet as an organizing tool generally did not occur until the later sessions. S1 began using the strategy as early as the 4th session, while S3 did not use the technique until the very last session.

S1's use of the alphabet as a generating tool is of particular note. She used the technique extensively for nearly 5 hours. Initially she covered the entire alphabet in less than 1 hour. She then made succeeding "deeper" passes using as many as four alphabetically generated letters to start a name. In the 8th session the entire hour was devoted to female names from Be-- to De--. The following is a sample of the protocol from S1's 8th session.

S1: O.K. Ummm. "C""E," Cecilia. Cecilia? Cecelia?
Oh! There was one younger. Umm. There was Camellia.
I already named Cemellia Blocker. Umm. Off this paper. [gesturing to a piece of paper which had the names of people the subject had recalled between sessions.] Let's see. "C." "C""E." Cecilia, umm, which can be like Cecil or Cil. Umm. "C""E." Umm. Celeste. Either as a first name or a last name. Celeste. Umm. No. Celeste. O.K. Umm. "C""E""M," "C""E""N," "C""E""P?" "C""E""R," "CERT." Umm. There's also Cary, "C""A""R""Y," but I didn't know any girls named Cary. Umm. "C""E""S," "C""E""T," "C""E""V," No. O.K. Umm. "C." "C""I?" "C""I?" Cinnamon? Umm. Silver? No that's "S." It'd be a last name. Umm. Cid?

5. This is also a good example of an overshoot and the value of a verification process to eliminate incorrect recalls.

It is interesting to note that this strategy which is employed comparatively late in the retrieval process is distinct from all the other strategies mentioned. The other strategies appear to have the goal of first identifying an individual and seeking that person's name. This strategy on the other hand, produces a name and then tries to find a person who has the name. It is as though the other strategies apply likely constraints to improve the chance that any name recovered will be a target name (e.g., activity to group to individual to name) while this strategy looks for constraints after the recovery of a target name (name to individual to group, i.e., "in my high school?")

Pictures. Another common strategy used by subjects was what I call the pictures strategy. Subjects would scan some internal "image" of a picture or set of pictures to locate individuals whom they would then attempt to name. Frequently the pictures were out of the subject's yearbook or from a collection which the subject had strong a priori reason to believe would have classmates pictured (e.g. a summer camp picnic). The following is a sample of this type of search:

S1: O.K. who else? Umm, Ronnie Walker, she was also the last letter. if I see her picture I can-- like there was a line-up of girls who were on the intramural volleyball team. Brenda was one. There were some younger people too. And Ronnie Walker was one... and... I already named Bret Hastings... Umm. There was another girl that was real good friends with Ronnie Walker. And she had long blond hair, and she was sort of tomboyish,... I can't think of her name. Hmm, let me see, if I go down anymore people in that picture...that I recognize. There was also the gymnastic picture. But most of those people were either younger or older. There weren't too many people my age level that was in that. Umm, I named Peter Walkoe, but I'm not sure if the other day, if had--that was like a if he had Walkoe then Walker came to me. Let me see. Umm, in the picture. Hmm...Jeff Peterson, I think I named him yesterday though.

The protocol sounds as if the subjects view veridical images. However, there is some evidence to suggest that these images can have interesting flaws in them. For example the following protocol is a sample of picture searching by S3:

I'm going through the yearbook again. I'm trying to look over the faces and maybe even try to visualize the names next to them, next to the pictures. Maybe that will help. I don't think it will. Sharon Farley. She was pictured next to Mark Farley and I remember her.

Sharon Farley.

It all sounds very plausible. And indeed, both Sharon Farley and Mark Farley are in S3's class. However, Mark Farley's picture is not in the yearbook! He is listed among the members of S3's class who did not have their pictures taken. The image from which S3 is abstracting her search information is wrong; it never existed in the real world.

Incidental Observations

In addition to the recurrent phenomena and search strategies listed above, the four subjects showed a number of characteristics which are better attributed to each of the subjects individually.

Some observations on S3. There are a number of aspects of S3's data which differ from the other three subjects. One is that S3 spent a larger portion of her time in a variety of verification processes such as overshoots and self corrections. Frequently candidate names had to be thought about for several minutes to insure that the names belonged to classmates in the correct high school. For example:

There was a girl named Bartholomew, but was she in Chula Vista or was she in Clairmont? Brenda Bartholomew, that was her name. There was this girl named Brenda Barthomew and I'm trying to find out if she was in Chula Vista or whether she was in Clairmont, and I think she was. Was she in . . . was she in Mr. McBrows class. God, I can't remember. She had very delicate features. Had very fine features on her face, small, she was sort of, not fat, but heavily built, and shorter than I was. Well, where was she. I can't remember whether she was in Chula Vista or Clairemont. I think in English class. I'm trying to place her in a seating arrangement in the classroom, and who was sitting around her.

From figure 3 we might note two additional things which distinguish S3's data from that of the other subjects. One, S3 recalled a much smaller percentage of the correct names possible. She recalled 93 out of 750 possible, for a percentage of 12.4, while S1, S2, and S4 recalled 35.1%, 29.6%, and 28.3% respectively. Two, there are a couple of dramatic jumps in the rate of names correctly recalled in S3's 4th and 11th sessions.

The difference in style may have been due to the fact that S3 went to more than one high school, and had lived in a number of locations within San Diego and other regions. S3 herself suggested at several points that it was difficult to distinguish

the people of one high school from those of the other.

S3: "B"'s.. Beutell, I already mentioned her. Who was around? I keep getting glimpses of Chula Vista now, Umm. Every time I think of "B," I think of Bolger, Brian Bolger, but that was Chula Vista. What else? Every time I think of the letter "B," I think of Chula Vista. I think a lot of people whose last names began with "B" were in Chula Vista, I can't get the Clairmont people straight. [several minutes later] Who else with the last name "S"? "S""A"? Sanford? That was Chula Vista, no, no. Let's get back to the "S"s again.

In other words, S3 may have been overwhelmed by distractions.

A second factor which may have contributed to the need for extensive verification was that many of the activities chosen by S3 as search contexts were situations where students from a variety of grade levels might be expected. For example, S3 frequently used the school band, language classes, and the international club as search contexts.

The two major cases of discontinuity on the recall of correct names in S3's data have some interesting characteristics. Both cases were preceded by a significant number of correct names recalled between sessions. At the start of the 4th session, S3 reported that in the interval between sessions she had accidentally met a person from her high school graduating class whom she had not named. Meeting that person immediately reminded her of a group of people she had not named. Throughout the rest of the 4th session S3 used the names recovered to suggest new contexts and to produce information which guided other searches. The 11th session had much of the same character. The subject mentioned recalling a new name which brought to mind a whole new set of people. The new names again appeared to spark a rejuvenation of the retrieval process. These discontinuities suggest a new line of experimentation. What happens when the experimenter introduces some new information to the subjects after the recall task has been running for some period of time? For example, what if the subjects were taken back physically to the high schools from which they graduated? Or what if subjects from the same high school were allowed to exchange information for some brief period?

Discussion

I interpret remembering as a problem solving process. To establish this point of view and to provide a theoretical framework for the interpretation of the memory search phenomena observed above it is necessary to analyze the task of retrieval.

I begin with what I believe are three primary constraints on the retrieval process: Partial Information, Descriptions, and Large Memory Capacity. ⁶ Partial Information: I believe that a person can encode only a limited amount of the possible information present in the environment. One can think of this limited amount of information as a list of features or properties, partial images, or whatever. The important point is that not everything that is present gets represented in memory. The partial information is a description of the actual event. Descriptions: A description is a theoretical retrieval term. To encode or retrieve any packet of information from memory a partial description is formed that provides an initial entry point into the memory. See the discussion of descriptions in memory retrieval by Norman and Bobrow (Note 3).

Large Memory Capacity. I assume that human memory can be treated as if it were indefinitely large, so large that there is always room for more information. In some sense, this is an anti-assumption. I wish to avoid any inclination toward thinking in terms of "economy of storage" or other similar notions.

From these notions, I can characterize a "kernel" retrieval process which uses some limited amount of information to form a description to access some information from memory. ⁷ From the location indicated, a record of information about the target item is retrieved. This record contains some limited amount of new information.

The "global" act of retrieval (i.e. what we observe in human behavior) can be characterized as a reconstructive process. The initial information is used to form a description of some aspect of the target item used in the kernel process. The information recovered in the kernel process is used in conjunction with what is already known to form a description, which is used to recover still more information. A succession of kernel retrievals are used to reconstruct the target item. Some of these ideas are well known. Norman (1968) discussed a number of retrieval problems including the problem of verifying the collection of the information, and the nature of the retrieval. Shiffrin (1970) examined some of these ideas in detail, and his work has influenced this characterization. Norman and Bobrow's (Note 3) characterization was developed, in part, from consideration of my data on very long term retrieval.

6. The notion of descriptions is taken from Bobrow and Norman, 1975, and Williams and Norman, (Note 2).

7. That some initial information or "retrieval specification" is necessary to initiate the retrieval process is clear if one attempts to imagine recalling anything without such a retrieval specification.

The kernel retrieval process of such a system runs into three basic problems: Too much information, too little information, and false recovery. 1. If there is more information at the time of retrieval than was available at the time of storage, then the retrieval specification may be over-specific, thus failing. Some mechanism must be found to reduce the information available to that which was used at the time of encoding. This is the problem of too much information. 2. If there is less information at the time of retrieval than was used at the time of storage, then the missing information must either be retrieved in some new retrieval cycle or it must be inferred. This is the problem of too little information. 3. Because partial descriptions are used to guide retrievals, a record similar to the one being sought may be recovered. I call this a false recovery.

The Retrieval Cycle

The retrieval cycle seems to be characterized by three phases (Figure 5), each intended to confront one of the basic problems above. First, there is the establishment of a retrieval context. This is designed to focus on a relevant subset of the initial information provided (thus, confronting the problem of too much information). There is a search through the memory space defined by the context and the available information (this is an attempt to combat the problem of too little information). Third, there is a verification of the information retrieved to minimize the chances of false recovery. If the information retrieved satisfies the original query of memory, the retrieval terminates at this point. Otherwise, the retrieved information is used to reformulate the description that guides the retrieval, and a new cycle is initiated. This new cycle may or may not require the establishment of a new context.

One important aspect of the retrieval cycle is that it is recursive. The establishment of a context, for example, may itself require a retrieval cycle, involving the finding of a context, the search, and the verification. Once the context is established, the search can begin, but it too may also require one or more retrieval cycles, this time contained within the memory established by the higher level establishment of a search context. Finally, the verification phase requires its own retrieval cycles, this time for the purpose of certifying the accuracy of the information provided by the preceding search phase. The interactive and related nature of the retrieval cycle was specified in detail by Shiffrin, 1970, although my characterization differs from his.

The Jigsaw Puzzle Metaphor

Imagine the problem of retrieval as being analogous to the reconstruction of a jigsaw puzzle when presented with the pieces from several similar puzzles mixed together. The information provided in the basic query specifies a number of starter pieces.




Figure 5. Basic characterization of the retrieval process.

We must use this information to find each new piece of the puzzle. We begin by restricting our view to some likely section of the puzzle, possibly where we have several pieces (locating this region might be thought of as analogous to finding a context). Next, we search for a likely piece by looking for something that will match a side with a piece in the "context." We rummage about in the search domain and pick out a piece that fits the description we have constructed (i.e., has a side which will complement a side of a piece from the context). If it were the case that too many pieces match the description, we might want to look for a bordering piece to get some more constraints on the piece we are seeking, or we might want to try a set of possible pieces (these actions might be thought of as conducting a search). Finally, we try to fit the piece we have selected to ensure that it also fits neighboring pieces, or we might use the newly recovered piece to locate new pieces (this is analogous to conducting a verification).

Contexts

The extensive use of search contexts by subjects bears directly upon the retrieval problem. At the time of encoding how does one choose a set of properties to use as a description? There are two competing factors: The properties should produce a description of a unique record. Otherwise, with a limited description, we might expect to have to search through a set of records to find the correct one. Secondly, the properties must be such that it is possible to generate the properties at the time of retrieval. The first factor suggests that we use as many and as diverse a set of properties as possible. The second factor suggests that we use some limited and stereotyped set of properties for encoding. The notion of context is a compromise between the two extremes.

A context is some procedure for selecting a specific subset of the properties which can be abstracted from any stimulus item. The properties should be "rich" enough that one can use a distinctive set for many different objects, yet they can be readily generated to facilitate the recovery of information stored using them as an index. Thus the context "my 10th grade art class" provides a specific set of characteristics which can be used to build descriptions to retrieve events which in turn lead to particular individuals.

Retrieval of a context. The context simplifies the problem of forming a description. However, the recovery of the appropriate context is itself a retrieval problem. Therefore the first subprocess of the retrieval cycle is the "retrieval of a context."

There are several ways one can imagine the retrieval of a context being accomplished. The most straight forward is to imagine that a unique supra-context exists for the single purpose of

finding good search contexts. This supra-context specifies what characteristics should be abstracted from the initial question. These characteristics are then used to build a description of the appropriate search context. Perhaps a more general way to retrieve a search context is with a recursive call to the retrieval process. In such a case a succession of contexts might be used to find the appropriate context to pass along to the search process. Thus, a subject might search the general context of high schools to retrieve a context of the football team which can then be searched for likely target items.

The importance of contexts, with regard to the retrieval process, is that there are far fewer contexts than items. Thus the retrieval of a context is easier than the retrieval of an item, and once an appropriate context is retrieved, it is easier to retrieve an item than it was before.

Search within a context. Given a context within which to work, the retrieval process will frequently be confronted with the problem of too little information. That is, the retrieval cues provided in a question may be insufficient (in and of themselves) to uniquely specify the record being sought. In such a case, a description based system has two alternatives. The first alternative is to recover records about the item (from some other context) for which a sufficient set of retrieval cues does exist. In this manner more information about the item is retrieved from which still more information can be gathered until enough retrieval cues are recovered to generate a good retrieval specification. The second alternative is to check out all of the records which can be retrieved with the limited set of cues available by hypothesizing about the missing cues. For example, recalling all but the last digit of a phone number, one could generate each of the 10 possibilities to see if one might be recognized. Note that as one generates each possibility one still has to check it out. That is, one has to attempt to verify the accuracy of the recovery in some manner.

The Recurrent Phenomena

The phenomena of extended retrieval, systematic hypothesizing, and inferential recall all are natural consequences of the problem of too little information. If less information is available at the time of retrieval than was used during storage, then additional information must be retrieved. The recurrent phenomenon I have called extended retrieval is explicitly the attempt to gather incidental information in order to enhance the description of the item sought. Two other recurrent phenomena, systematic hypothesizing and inferential recall are both examples of the enormous amount of work a subject will go to in order to recover a necessary fragment of information. In both cases the subject is taking a chance of making what is essentially an

intentional false recovery⁸ in order to continue a search that is working with too little information, in the hope that new information will be recovered and can be verified.

The phenomena of overshoot, fabrications, and self correction are evidence of various levels of the problem of false recoveries. Overshoot and self correction can be interpreted as direct evidence of verification processes which are necessitated by false recoveries. (The problem of verification and its implications are discussed in greater detail in the next section.)

The Interpretation of Some Memory Phenomena

The notions of partial information and descriptions provides a framework in which a number of traditional memory phenomena as well as those observed in the protocols above can be interpreted.

Forgetting. One of the most salient phenomena of human memory is forgetting. Frequently memory models account for this phenomenon by introducing notions of decay or unlearning. The characterization of memory I have introduced has no provisions for the decay or deletion of information once it has been stored. How then is it possible for this characterization to account for forgetting? I believe that forgetting is caused by four problems which confront the retrieval process. Those problems are: false recovery, re-encoding, too little information, and too much information. As a result of these problems a variety of conditions can make a specific piece of information more difficult to recover. If the difficulty becomes too great, the search is abandoned and the information is reported to be "forgotten."⁹

1. False recovery. Any time an event is being reconstructed (i.e. recalled), it is possible to recover a record about a similar event and incorporate it into the image of the event being recalled. When this "merging" of two events occurs a new event has been constructed. A false recovery which is accepted creates a fabrication. Because the retrieval process is using information about the event to find contexts, to search within contexts, and to verify the recoveries, the retrieval of new information

8. I define a false recovery as the retrieval of information about an item which was not explicitly encoded about that item. Therefore a false recovery could be the recall of a piece of information about some similar item, or it could be the recovery of a piece of information about some more generic item or mental event.

9. Notions similar to this have been proposed before in the literature. Tulving and Pearlstone, 1966, with their distinction between availability and accessibility, are an example of such an idea. One important difference is that what I am proposing is a mechanism for such retrieval failures.

can become blocked. The fabricated event never actually occurred, therefore no information about it exists.

Among other things, the problem of false recovery suggests that the more events of a similar nature that have been encoded in memory, the more difficult the task of retrieval becomes.¹⁰

2. Re-encoding. Each time we recall the information we have stored about an object or event we recall less than was initially encoded. This happens for several reasons. For one, generally we need only some limited amount of information about the item to perform the task at hand. Therefore we can quit recovering information once we have enough. For two, frequently some of the information is forgotten for reasons discussed in the other paragraphs of this section.

Any time we recall information about an event we have an opportunity to re-encode it. Accordingly, when we attempt to recover information about the event the next time, we might recall the original event or the re-encoding. The iteration of this process results in an object or event being re-encoded with less and less information.¹¹

The problem of false recovery suggests that in addition to being less complete, re-encodings may be fabrications. This would accelerate and compound the forgetting due to re-encoding.

When recalling an event we may fill in information which has not been explicitly recovered (perhaps because it was never encoded, or perhaps because it has become too difficult to retrieve). In addition to filling in bits and pieces of the information about an event, we may also reprocess the information we have recovered to discover new things.

Additions to the re-encoding of an event may make access to the original encoding even more difficult. The new pieces of information and the new interpretations are not marked. If we look for the original encoding using this new information or in contexts suggested by these new interpretations, we must fail. Re-encodings may also improve the ability to recover the information which was in the original encoding and not in the re-encoded version of the event.

10. Thus, response competition, retro/pro-active inhibition, and other like phenomena are a consequence of this mechanism of retrieval failure.

11. Such a course of events might well lead to the distinction between semantic and episodic memory that Tulving (1972) has propounded. Bartlett's (1938) observation that repeated recalls becoming successively more stereotyped also fits well with this mechanism.

It is interesting to note that mnemonists report that when information is encoded using a mnemonic strategy, after the information has been recovered a number of times using the mnemonic scheme, that the mnemonic seems to disappear and all that is left is the information which was originally encoded (see Lorayne and Lucas, 1974).

3. Too little information. Too little information is the case where less information is available at the time of retrieval than was used during storage. In this condition the retrieval process confronts three problems. One, a false recovery is much more likely. This is because little information is available to use in the verification process and because the number of possible events described by the information is large. Two, it becomes more likely to recall a re-encoding of the event than the original event. This occurs because re-encodings are encoded using fewer and perhaps more typical properties. Three, if a property which was used for indexing part of the knowledge about an event is not accessible from the information presented in the question, then the knowledge indexed under that property can only be recovered by guessing. If a few such properties are not available, the guessing combinations become numerous.

4. Too much information. Too much information is the condition where more information is available at the time of retrieval than was present during storage. In this case the retrieval problem is essentially one of trying to determine the correct context to use. If the subject is misled into choosing the wrong context, the retrieval task becomes impossible.

All of the retrieval problems talked about in this section interact with one another creating difficulties for the retrieval process resulting in "forgetting." The problems of too little information and too much information may occur on different cycles of the same retrieval task. Subjects may begin a task with too much information. Once they find the correct context, they may not have enough information within that context. Too little information is a problem, in part, because of the increased probability of a false recovery, or of recalling a re-encoding of the item sought.

Search. The retrieval cycle is an iterative process. Information is used to construct a description, which is used to locate a record. The information recovered is used in conjunction with what is already known to recover still more information. As more and more information is recovered the target is narrowed down until a small set of fragments remain. These fragments are reconstructed into an interlocking puzzle or image of the item from which the target information is abstracted. The phenomenon of search is the natural product of this process.

On the human scale, search is a common enough phenomenon. Any difficult retrieval task, almost by definition, results in a search. What is surprising, from the point of view of my characterization of retrieval, is the occasional "instantaneous" recollection. Sometimes a recall appears to take no search at all.

I suggest that even these rapid retrievals require some search; every retrieval subject to experimentation involves a search. The difference lies in the length and the complexity of the searches. The common phenomenon of search is the result of those retrievals which take long enough for the intermediate steps of the retrieval to be distinguished. Even in the recalls that appear to be instantaneous, information beyond what is sought frequently comes to mind. Whether this other information occurred before or after the retrieval of the target information is uncertain. My view is that all recollection can be placed along a dimension of the amount of search. In lengthy searches the search processes can be easily identified, while recollections which appear to be instantaneous are simply brief searches.

There are a variety of phenomena which can be interpreted as evidence for this view that even "instantaneous" or immediate recalls are the product of reconstructive search processes. In brief, they are: 1. The continuity of recognition reaction times. If I delay the recognition process, the time required to recognize an item increases. Thus if I test subjects immediately after the presentation of a list, they will recognize items faster than if I test them after a delay (Woodworth, 1938, pg.38). This increase in reaction times can span from immediate recognitions with no apparent search to lengthy recalls in which search protocols can be obtained (Boeck, Note 4). There is no sharp increase in reaction times which might suggest distinct processes.¹² 2. Intrusions. The errors which occur in instantaneous recalls have the same reconstructive character as those errors obtained from lengthy searches. A prime example would be speech errors. The retrieval time per word is very brief. Nonetheless, speech errors often have the character of being an item which is similar to the target item on any of a variety of dimensions, or the improper reconstruction of the correct pieces. 3. Retrospective reporting. Occasionally, subjects can report an involved sequence of retrieval steps which lead to a recall which occurred in a brief instant. For example, the following is a segment of S1's protocol during an extended period of a Location search strategy:

¹². Though Waugh (1970) has demonstrated that if the delay interval is brief enough to allow the presence of the item in primary, as opposed to secondary, memory that reaction times are substantially reduced.

S1: Okay, let me see. Does anybody else live back in there that I know now? Oh! Maxine Levin. She didn't live there, but Pam Volpone now lives on Loma Portal and Mike Stevenson, who is Eleventh grade, used to live in the house that Pam Volpone and her husband bought. And so from Pam Volpone I got Maxine Levin 'cause I saw her at the store the other day and she was talking about Pam Volpone.

This protocol suggests that the subject first recovered a house in which a friend who was a year younger, Mike Stevenson, lived in; and from that recalled that a classmate, Pam Volpone [named earlier], now living in the house; and from that recalled that the first classmate, Pam Volpone, had been the recent subject of a discussion between the subject and another classmate, Maxine Levin. If the subject's report is correct, then all of these recollections occurred within the fraction of a second pause between sentences. 4. Tip of the tongue phenomenon. This phenomenon occurs naturally in normal discourse, and can be interpreted within our framework as a search failure resulting in extended retrievals, partial recalls, etc.

Recall and Recognition. In the characterization of memory presented in this paper, the common distinction between recall and recognition fades into the dimension of search. In the classic recall paradigm the subject is given a minimum of cues from which to begin a retrieval, while in the recognition paradigm the subject is given an abundance of cues from which to initiate the retrieval. The result in the first condition is a comparatively difficult and lengthy search, while in the second condition the search is generally easy and therefore brief. Mandler (Note 5) points out the frailty of a sharp distinction between recall and recognition, and the common confusion between the theoretical notions, the experimental paradigms, and the colloquial meanings of the two terms.

It is my view that it is not useful to make a sharp distinction between recall and recognition. Rather, one should think of the two notions as two poles along a dimension of density of useful retrieval information. As the amount of information presented increases, remembering becomes more like recognition and less like recall. While the problem in most recall paradigms is too little information (resulting in the search for more information), the problem in most recognition paradigms is too much information (resulting in the need to choose the appropriate context).

Intrusions. Intrusions are errors in the recollection of an item. Frequently they are considered to be any kind of erroneous word recalled in a standard word learning experiment. Sometimes the term can refer to the recollection of an event which is the confounding of two or more separate events. Intrusions have also

been referred to as fabrications, or simply errors.

The characterization of the retrieval process which I have developed here has as its logical consequences not only the simple fact of intrusions but also many of their observable characteristics.

From the point of view of my characterization of memory intrusions, come both from false recoveries and from inferential recalls. Because only partial information can be used in the specification of a record, the record which is recovered may be the one originally encoded about the event being sought, or it may be a record about a similar yet distinct event. This false recovery may contain information which is not true with respect to the event being sought. These false recoveries have several results:

1. The false recovery can be incorporated in the reconstruction of the information being output. For example, the acoustic errors which were observed as fabrications.

2. The false recovery can misdirect the search for more information. This misdirection can itself have several results.

- 2a. The false recovery can cause the search to fail. The information recovered may not lead to proper records, and the false information may lead to faulty verification, resulting in the rejection of appropriate records or directions for further search. The following protocol from S1 is an example of a search being misdirected by a false recovery.

Nancy Phillipano, is that other Nancy. I couldn't think of her name, and it really bugged me. She was a girl that was on ROTC. This is a different Nancy all together. Uhhh... She was a girl that was on ROTC, and she was sort of dingy, and she had a sister that was a year younger, but also graduated with our class, and her name was. Because she would have been in my senior class too. Last name was Phillipano. There was Nancy Phillipano and her sister, who was.... Her name began with an "R." Rochelle, or, ummm, not Rochelle. Her sister was sort of pretty, and she always wore it, it was long and sort of real real curly, and she'd always wear it like with the top part pulled back into a pony tail in the back and would let the back hang down. She would always wear a suede jacket. And what was her name? Seems to me that it began with an "R." ummm, let me see... Not Robbie, not Rachel, ummm... Phillipano, Nancy Phillipano. And what was her sister's name? Not Rachel, ummm... Not Ramona, not ummm, let me see. O.K. I'm trying to think of "R" names. She had a sort of unusual, not Rachel, not Roberta, not Robin, Robin, Robin Phillipano? No. Roberta Phillipano? No. Ummm... Let me see. Ummm, let me see, ummm. Not Rolly. Ummm. Can't think of any more "R" names. maybe (mumble) With an "R." But

she did graduate with our class. She graduated a year early. Her last name was Phillipano. Who did she hang around with? Let me see, of other people who are (mumble). Diane Hart! That's a girl I couldn't think of her name. That's one girl's name. That's a new name.

In another session S1 recalls - correctly - that Nancy Phillipano's sister is named Luann. As long as S1 is looking for a name which begins with an "R," Luann cannot be recovered.

2b. The search may be directed to additional false recoveries, resulting in the fabrication of an event which never actually occurred, or in the recovery of an item similar to the target item. For example, the misassignment of names which were observed as fabrications.

Inferential recalls could be considered intentional false recoveries. Inferential recalls are recollections about items similar to the one sought about the general class of items of which the target item is a member. Thus, in recalling that a person's name is "German sounding," the subject goes on to recall information about "German sounding names" to be used in guiding the search. Each "German sounding name" that is recalled in this fashion is, technically, an intrusion.

Inferential recalls can have the same effects on the retrieval process as do false recoveries. Thus, an inference can become a part (or a whole) of an output, and can occasionally misdirect the search. Note that inferential recalls need not necessarily misdirect the search; indeed, their main function is to provide intelligent guidance to the search process.

Verification. The verification process is a natural consequence of the problem of false recoveries and the necessity to confirm inferential recalls. The verification process is meant to identify false recoveries and to reduce the possibility of being misled by an inferential recall.

In the retrieval system I have described, there are several ways information can be verified. One technique is coincident recovery. Once a piece of information has been recovered it can be considered as simply a possibility until it has been confirmed by an independent recovery of the same information. A second verification technique is indirect confirmation. If a subject can use a bit of information just recovered as part of a description to retrieve additional information, then the subject could judge the information to be verified. This technique may be the source of the overshoots observed. A third possible verification technique is consistency checking. If the information recovered fits with what is already known then a subject could consider the information likely to be correct. Thus if a subject recalls that a particular person who is known to be on the football team is big and heavy set, then this new information is consistent with

what is already known about football teams in general.

Isolating verification techniques in this manner, it is possible to think of the verification process as a set of sub-processes, each of which filters out some percentage of errors.

¹³ In practice it is probably the case that all of these techniques (and others not considered) are used together in various combinations during any specific recall. Thus applying various filters to reduce, but not eliminate, retrieval errors.

Another way of thinking about the verification process is as an integration of the application of the techniques mentioned. Thus a subject might judge the "connectedness" or consistency of a large array of fragments of information recovered about a particular person. Judgements of how tightly bound any piece of information is to the whole (e.g.: Is it consistent with what is known?; Can it be used to recover other information which is known?; Can it be recovered in more than one way?) may serve as the basis for judgements of certainty. Though this second way of conceiving of the verification process is more complicated, it is the view I favor.

I propose that the degree of certainty that subjects express about the validity of their recollection is based upon the extent of success they have had in verifying them. If the recollection hangs together well, if each piece of information recovered is consistent with each other piece, if additional information is readily accessed using the information which has already been recovered, if the information to be output can be recovered in more than one way, then the subject will express a great deal of confidence in the accuracy of the output. If there are anomalies in the recollection, if no additional information is available using what has already been recovered as a tool for searching, if the information that has been recalled can only be recovered in one manner, then the subject will express less confidence in the output.

Distractions. Distractions are pieces of information that subjects recover which they know to be false recoveries, but which, nevertheless, recur and interfere with correct recalls. An example of a distraction and a more detailed explication of the phenomenon is presented in the section on recurrent phenomena. Briefly, the three essential characteristics of distractions are: one, they interfere with the retrieval process, two, subjects like to pinpoint the source of distractions (some claiming this aides in reducing the distractions interference), three, distractions match the target item only on a limited domain.

13. These filters also might eliminate good information along with the bad.

Distractions can be interpreted as the consequence of my characterization of the retrieval system. The retrieval of a distraction is the same as any fabrication except that the subject realizes that the recovery is an error. Identifying the source of the distraction permits the creation of a set of criteria for discriminating (during the verification process) the records recovered about the target item from those of the distraction. Suppose I know two Larrys, Larry1 and Larry2, with the desired target being Larry1. If I recover the characteristics of Larry2 (Distraction), (e.g. UCSD professor, last name Shepard, etc.), then I have a means of discriminating records about Larry1 from Larry2.

The fact that distractions which match the target item on only a limited set of characteristics still intrude is indicative of the use of search contexts. That Larry1 and Larry2 differ on the characteristic of where the subject knew them has no impact on eliminating the distraction from initial recovery.¹⁴ That is because location is not a part of the NAME context, while "Larry" is.¹⁵

Conclusion

I have attempted to do two things: one, present a collection of my observations from verbal protocols produced by subjects thinking out loud while recalling the names of their classmates in high school, and two, introduce the framework from which to build a theory of retrieval from long term memory which provides an integrated explanation of what I have observed as well as an array of traditional memory phenomena.

My observations were taken from four subjects who were thinking out loud while attempting to recall the names of their high school classmates. The subjects, who were from 4 to 19 years out of high school, engaged in this task for brief, approximately one hour, sessions for total recall times of from 4 to 10 hours. Observations included a variety of recurrent phenomena exhibited by every subject, a set of common search strategies, and a collection of incidental phenomena which appeared to me to be interesting though difficult to classify.

The recurrent phenomena observed included: Partial recalls, Incidental recalls, Search Contexts, Extended Retrievals, Inferential Recalls, Overshoots, Fabrications or intrusions, Self Corrections, and Distractions.

14. Though it does have an impact on eliminating records during the verification stage.

15. This observation, if correct, points out some of the limitations of a simple set theoretic search models which do not employ contexts, or some other mechanisms, to focus the search effort.

I also observed the enormous amount of information that subjects can recall if they are pressed. Though the subjects frequently commented during the initial sessions of the task that they didn't think they would be able to recall any additional names, they were able to go on recalling names virtually indefinitely.

The common search strategies that subjects used included: General Association, Activity search, Location scanning, Name Generation, and Image scanning.

Though many of these strategies are probably peculiar to the specific task given to these subjects, the fact of the clear presence of these strategies and the individual differences in the depth and duration of their application are characteristics of the search process which should be represented in any serious description of long term memory retrieval. ¹⁶

The theoretical interpretation introduced is based on the view that retrieval is largely a problem solving process. From the three basic constraints which I have called partial information, descriptions, and large memory capacity I have suggest a three stage characterization of the retrieval process. This characterization was used to provide an interpretation of a variety of classical memory phenomena as well as many of the recurrent phenomena mentioned in this paper.

The characterization of the search process is substantially that of a reconstructive retrieval process. Information about the target item is used to construct a description of some aspect of the item. The description is used to recover a fragment of information about the item which is added to what is known. From this information a new description can be formed to retrieve still more information, until the particular piece of information sought can be recovered. The three stages are FIND A CONTEXT, in which a proper environment for conducting a search is recovered, SEARCH, in which bits and pieces of information appropriate to the context are recovered until an adequate description can be formed within the search context, and VERIFY, in which the record recovered is checked to confirm that it is about the target item. Each of the three stages has embedded within one or more recursive calls to the retrieval process. This characterization was used to interpret the recurrent phenomena observed in subject protocols, as well as memory phenomena such as forgetting, search, recall and recognition, and intrusions.

16. Work taken in this direction which has recently come to my attention is that of Keniston and Flavell (note 1).

Reference Notes

1. Keniston, A., and Flavell, J.H. The nature and development of intelligent retrieval. Unpublished manuscript; 1978.
2. Williams, M.D., and Norman, D.A. Some characteristics of retrieval from very long term memory, paper presented at the meeting of the Psychonomics Society, St. Louis, MO, 1976.
3. Norman, D.A. and Bobrow, D. Descriptions: A basis for memory acquisition and retrieval. (CHIP report No. 74) La Jolla, University of California at San Diego, Center for Human Information Processing.
4. Boeck, W.J. Personal communication, October 1976. Reporting some pilot results which are being expanded in his Ph.D. dissertation at the University of California, San Diego.
5. Mandler, G.M. Memory research reconsidered. (CHIP report No. 67) La Jolla, University of California at San Diego, Center for Human Information Processing.

References

- Bartlett, F. C. Remembering. Cambridge: Cambridge University Press, 1932.
- Bernbach, H.A. A multiple-copy model for postperceptual memory. In D.A. Norman (Ed.), Models of Memory. New York: Academic Press, 1970.
- Bobrow, D. G., & Norman, D. A. Some principles of memory schemata. In D. G. Bobrow & A. M. Collins (Eds.), Representation and understanding: Studies in Cognitive Science. New York: Academic Press, 1975.
- Erdelyi, M.H., and Kleinbard, J. Has Ebbinghaus decayed with time?: the growth of recall (hypernesia) over days. JEP: Human Learning and Memory, 1978, 4, 275-289.
- Nisbett, R.E., and Wilson, T.D. Telling more than we can know: verbal reports on mental processes. Psychological Review. 1977, 84, 231,259.
- Norman, D. A. Toward a theory of memory and attention. Psychological Review, 1968, 75, 522-536.
- Shiffrin, R. M. Memory search. In D. A. Norman (Ed.), Models of human memory. New York: Academic Press, 1970.
- Tulving, E. Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory. New York: Academic Press, 1972.
- Woodworth, R. S. Experimental psychology. New York: Holt, 1938.