Handbook
of Learning and
Cognitive Processes

Volume 1

Introduction to
Concepts and Issues

EDITED BY

W. K. ESTES

Rockefeller University

Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review,* 1956, **63,** 81–97.

Miller, N. E. Applications of learning and biofeedback to psychiatry and medicine. In A. M. Freedman, H. I. Kaplan, & B. J. Sadock (Eds.), *Comprehensive textbook of psychiatry.* (2nd ed.) Baltimore, Maryland: Williams & Wilkins, 1974.

Murdock, B. B., Jr. *Human memory: Theory and data.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1974.

Newell, A., & Simon, H. A. GPS, a program that simulates human thought. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought.* New York: McGraw-Hill, 1963.

Pavlov, I. P. *Conditional reflexes.* (Translated by G. V. Anrep) London: Oxford University Press, 1927. (Reprinted by Dover, New York, 1960.)

Sears, R. R. Survey of objective studies of psychoanalytic concepts. A report prepared for the Committee on Social Adjustment. *Social Science Research Council Bulletin* **51,** 1943.

Shannon, C. E., & Weaver, W. *The mathematical theory of communication.* Urbana, Illinois: University of Illinois Press, 1949.

Skinner, B. F. *The behavior of organisms.* Appleton-Century-Crofts, 1938.

Suppes, P. Stimulus–response theory of automata and TOTE hierarchies: A reply to Arbib. *Psychological Review,* 1969, **76,** 511–514.

Thorndike, E. L. *Human learning.* New York: Century, 1931.

Titchener, E. B. *Experimental psychology. Qualitative.* New York: Macmillan, 1909.

Tolman, E. C. *Purposive behavior in animals and men.* New York: Appleton-Century-Crofts, 1932.

Watson, J. B. *Behaviorism.* New York: W. W. Norton, 1924.

White, S. H. Evidence for a hierarchical arrangement of learning processes. In L. P. Lipsitt & C. C. Spiker (Eds.), *Advances in child development and behavior.* New York: Academic Press, 1965. Pp. 187–220.

Winograd, T. Understanding natural language. *Cognitive Psychology,* 1972, **3,** 1–72.

2

Cognitive Psychology: An Introduction

Gordon H. Bower

Stanford University

I.   INTRODUCTION

Cognitive psychology is concerned with how organisms cognize or gain knowledge about their world, and how they use that knowledge to guide decisions and perform effective actions. Knowledge enables us to survive in a hostile environment, to satisfy our social and biological needs, to plan for our own and our children's future. Organisms exhibit several broad classes of knowledge: knowledge that certain propositions are true (so-called factual knowledge), and knowledge of how to do a variety of things (skills, procedures). Historically, stimulus–response psychology emphasized motor skills, whereas cognitive psychology emphasized factual knowledge. But obviously an organism needs both types of knowledge to get along in this world. Because the cognitive system has developed as an instrument of adjustment, as an aid to satisfying needs and motives, it is appropriate that we begin with discussion of biological and social motives, and their directive function.

A.   Motives Driving the Cognitive System

1.   Biological Motives

We have evolved as biological machines under the selective pressures of the "contingencies of survival," as Skinner (1974) calls them. Evolution eliminated those species that did not have adaptive exchanges with their environments. Body cells require periodic replenishing of their biochemical

nutrients; organisms that interact effectively with their environment to get those nutrients are the ones that survive.

A set of interconnected homeostatic mechanisms has evolved to handle the biological emergencies arising when the equilibrium of the body's internal milieu is disturbed. For example, when a person becomes hot, this condition is detected by heat-sensitive cells in the vascular bed of the hypothalamus. The input to those sensors passes some threshold which switches on several vasomotor reflexes that will produce bodily cooling: the capillaries dilate, carrying more red blood to the skin surface for cooling; the skin sweats, so that evaporation of the water produces a cooling effect; and the organism is impelled (motivated) to carry out actions whose ultimate goal is cooling—he sheds excess clothing, he stands in a breeze or in the shade, he seeks an air-conditioned room, or he may say, "Would you please open the window?" etc. All such activities are instrumental means to the end of reestablishing homeostasis in the internal milieu.

Knowledge and motor skills play major roles in this adaptive process. First, as a precondition, the adaptive organism must be able to identify internal states (for example, cooling of the body core) that comprise favorable changes in the internal milieu. Second, through associative learning similar to Pavlovian conditioning, he must learn to identify external situations or happenings in the world which will bring about (or are at least correlated with) favorable internal changes. These situations (for example, a cool room) are subgoals, and they come to be valued as a means to a desired end. Third, the organism must be able to acquire and then apply coordinated action patterns that bring about these valued goals. This is the "situation–action–outcome" knowledge acquired through the process we call operant conditioning. For example, a person must learn how to turn on an air conditioner, or how to get to swimming pools and air conditioned rooms that are located nearly. It is particularly in regard to these latter learnings—in effective interactions with the environment—that the cognitive system seems to have developed to the highest degree.

2. Social Motives

We have illustrated the motives underlying behavior in terms of a biological need (temperature regulation), although other biological needs like hunger and thirst would have served as equally cogent illustrations. However, it is clear that organisms pursue many autonomous goals that are not of this impelling metabolic kind. We engage in many activities for their own sake—hiking, painting, building toys, playing music, playing games of all sorts, and so forth. We probably also have "cognitive motives," such

as the desire for conceptual clarity in puzzling situations, the desire to reduce uncertainty about future events, and the desire to know how or why things work. Maslow (1967) proposed that people have an ordered scale of motives, ranging from biological needs at the lowest end of the scale through social needs (such as affiliation, love) to more transcendent personal needs such as self-awareness, actualization, and personal growth at the highest end. The customs of modern societies are partly organized around satisfying automatically the lower motives. Maslow hypothesized that people would pursue motives at a given level only if those at lower levels had been satisfied. Judging from the evidence of social and personal disorganization in concentration camps, Maslow's conjecture is at least plausible.

3. Dealing with Multiple Motives

People are obviously guided by multiple motives. Several methods have envolved to deal with this multiplicity (see Simon, 1967). The first method is by "priority queueing" of goals in terms of their urgency: if, while I am pursuing satisfaction of Motive 1 (say, I am a salesman selling a car), conditions arise that give prominence to Motive 2 (say, going to the toilet), then depending on the relative urgencies the cognitive system can either (1) queue up the newly arrived goal in a waiting line to be serviced, or (2) interrupt activity towards Goal 1 in order to turn momentarily to Goal 2, and then return to Goal 1.

A final method available for dealing with multiple goals is to seek out activities which will satisfy several goals simultaneously. An example of a multigoal activity is the "business lunch": it serves the goal of doing business as well as reducing hunger. Furthermore, the exact form, surroundings, and contents of the lunch may satisfy other goals of aesthetics, quiet relaxation, and perhaps enhancement of self-esteem. Sometimes the main factor determining selection among two actions is the greater number of motives that can be satisfied by one rather than the other action.

B. Intentions, Purposes, and Actions

These matters of motives are mentioned at the outset to reveal a fundamental assumption of cognitive psychology, namely, that the unit of behavior is the "purposive action" rather than the "colorless movement" or "glandular squirt" of yesterday's behaviorists. In this regard, cognitive psychologists side with the layman who uses an action terminology committed to the idea that goals explain or account for the behavior. Thus, when asked "*What* is John doing?," we can describe the action by some

remark like, "He's trying to build a fire." But the very same answer also serves as the layman's *explanation* when we ask, "*Why* is John doing that?" In characterizing the action, we are also characterizing the intention or goal of the action, and the layman's theory explains behavior as an expression of intentions, desires, and the like.

Tolman (1932) and other cognitive psychologists thought such explanations were proper though incomplete; he recognized that one had yet to explain why exactly *this* intention had arisen and *this* class of actions had been chosen to achieve the goal—but Tolman supplied ready answers in terms of demands and past learning experiences. The stimulus–response behaviorists such as Watson, Guthrie, and Hull objected because they thought psychology should deal with objective movements—that is, muscle twitches and glandular secretions—and only explain purposive actions in terms of these elements. They thought explanations couched in terms of purposes and goals were teleological and necessarily unscientific.

Extensive philosophical analysis has been devoted to the concepts of actions, movements, intentions, and teleological explanations, (see particularly Taylor, 1964). The upshot, I think, is that the behaviorist's insistence upon colorless, purposeless movements as the "units" of behavior is best viewed as a mistake. Actions are simply not reducible to, or equivalent to, a set of movements. It is practically impossible to describe a response without reference to its achievements, or to the stimuli with which the effectors (muscle groups) are in contact, and to the feedback provided by these effectors. Moreover, a given movement sequence can be two different "actions" depending upon the inferred intention or "meaning" of the action to the actor. This difference is captured in our everyday distinction between "accidental" and "intentional" responses and consequences. If a man hits me in the face as he is stretching, I infer different intentions if I know he is asleep than if I know he can see exactly what he is doing. The distinction also appears throughout the law: for example, a defendant is charged with first-degree murder if the killing was intentional, is charged with manslaughter if the killing was unintentional but related to his "negligent" actions, but is not charged at all if the killing was clearly "accidental."

Finally, action units simplify descriptions of behavior since we aggregate behaviors according to their immanent effect or goal, and this suppresses large amounts of irrelevant variability at the molecular level of movements. Thus, for most theoretical purposes, it is adequate to say that a man is "building a fire," even though it is understood that there are many different ways to build fires depending upon circumstances, each with ideosyncratic variations. Action terms serve as classificatory concepts; because of this, it would not be inappropriate to say that the man is behaving in such manner as to produce or realize an instance of the action concept we call "fire

building." Our procedures for assigning a stretch of behavior to an action class are similar to those by which we assign physical objects and events to classes. A significant difference is that in order to qualify as an instance of action class X, the behavior must vary according to the circumstances in such manner as to achieve X in those circumstances. Another illustration of the open-ended character of action concepts arises with locomotion in space, as when we say, "I'm going to the grocery store." In getting from point A to point B in town, we know and could use multiple routes and means of conveyance as circumstances require: we could go in a wheelchair or crawl or hobble on stilts if required to, even though we have never gone from A to B in any of these ways. This is not to say that once I learn one way to go from A to B, I automatically know all possible ways to get from A to B (which is absurd). Some earlier sensorimotor learning is clearly needed in order to apply or transfer particular skills to the present problem.

In cognitive psychology, the intentional actions of a person are characterized in terms of his carrying out a *plan* which has a particular goal. That is, once a given motive or goal assumes top priority, the person selects a behavioral plan from memory that, given the present circumstances, should bring about that goal.

C. Plans and Hierarchies

1. Plans

This concept of a plan was introduced by Miller, Galanter, and Pribram (1960) and it appears throughout the writings of later psychologists. A plan may be thought of as a procedure or recipe for achieving some goal. Most plans have a hierarchical structure, which proves useful for analyzing large segments of goal-directed behavior. A given goal will generate several subordinate goals (called subgoals) possibly with their own subsubgoals. The logical arrangement of subgoals within goals can be diagramed as in Figure 1, which illustrates a plan that goes two levels deep. Suppose your top goal is to go to a theater play tonight. To do this, you must first reserve tickets (goal A). To reserve tickets, you must first look up the telephone number of the theater (do C), then telephone and reserve tickets (do D). If that succeeds, then this evening you must get to your seat at the theater on time (goal B). This requires leaving on time, driving your car to the theater, and parking (do E), then paying for your tickets and walking to your seat for the start of the play (do F). By satisfying subsubgoal F, you also satisfy subgoal B, which is a precondition for your satisfying
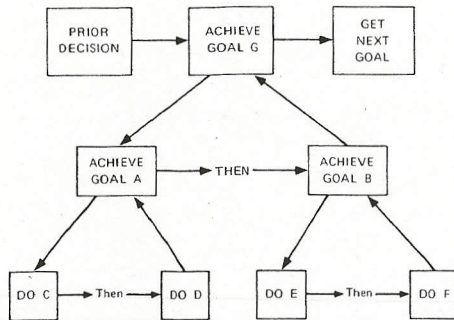
FIGURE 1 Diagram of a hierarchical arrangement of subgoals embedded within a top-level goal.

the top goal *G*. After that, you might begin operating on a new goal, such as finding a coffeehouse for refreshments after the concert.

These kinds of hierarchical analyses of action plans have considerable appeal; they are a neat way of parsing or chunking the significant aspects of behavior. One can refine and expand the behavioral description to as much detail as is of interest to the current investigation. Thus, the subgoal "look up telephone number" is a primitive in the description of Figure 1, but is differentiable in terms of other subgoals and elementary situation–action routines, for example, "find the phonebook," "open to desired page," "scan list for name." Each of these in turn is further differentiable, for example, "scan list" means to look at each entry successively until you find one that matches the name desired.

An attractive feature of this approach is that it solves the ancient "molar versus molecular" dispute concerning the proper level at which behavior should be characterized by the observing psychologist. As is well known, a given bit of behavior can be described at various levels; we could say that our subject is making particular subvocal movements, that he is comparing the name "Bijou" to his target, that he is moving his finger down the page, that he is looking up a phone number, that he is making ticket reservations, that he is planning to go to the theatre, that he is promoting his personal growth, and so on. It is senseless to ask which of these things is he "really" doing; the answer is that he is doing all of them, more or less at the same time. They are simply descriptions that focus on subgoals at different levels of the hierarchical plan he is carrying out.

These hierarchical descriptions of plans are quite congenial to the information-processing approach since plans resemble flow charts showing the passage of control among different components (the "subroutines") of a hierarchical program that is solving a problem. The behavioral plan may thus be represented by us as a computer program, a sequence of instructions and branching decisions for carrying out some actions that are likely to achieve some goal. When the person "executes" a plan, he proceeds through it step by step, completing one part and then moving on to the

next. A very elementary plan for "walking the length of a block" (Simon, 1967, p. 31) would be:

Walk the length of a block:

1. Step with left foot, then 2.
2. Step with right foot, then 3.
3. If end of block, do 4; if not, do 1.
4. Terminate.

This is a list of elementary instructions to be performed in order, with a conditional test inserted at the third line. A similar small program might be written for "how to cross an intersection." By combining these two programs, we can synthesize somewhat larger or more interesting segments of behavior, such as "Walk seven blocks north," thus:

Walk 7 blocks north:

1. Face north, then 2.
2. Set count = 0, then 3.
3. Walk the length of a block, then 4.
4. Add 1 to the count, then 5.
5. If count = 7, do 7; else, do 6.
6. Cross intersection, then 3.
7. Terminate.

The important point is that "walk a block" and "cross intersection," although themselves programs, can be used as units within a larger program. These smaller segments are called subroutines and may be characterized or named in terms of what functions they perform. In today's programs to simulate learning or problem solving, there would be a hierarchy of subroutines that are organized in some way designed to achieve particular goals.

You might ask: *What* reads and executes these alleged instructions in your alleged plans? In the computer, of course, it is the *interpreter* which "reads" instructions and translates them into physical happenings. But when one proposes a program (or mechanical algorithm) to simulate behavior, he need not be proposing that there is some inner interpreter (homunculus) that "reads" the physiological counterpart of the instructions and fires off the effector movements. Rather, the program is just *our description* of the system; it is external to the subject himself. The program need not correspond to anything inside the person; rather, the program may only behave the same way as the mechanisms that do control the behavior. For instance, one could write a program to simulate a thermostat, but he would not thereby be committed to the assertion that the expansion wire in the thermostat is "following" that program. In contrast, there may

be in a real system something that does correspond to a program that causally controls its behavior. An example would be the pattern of holes in the paper player roll that controls the movement of keys in an old-fashioned player piano (the examples come from Newell & Simon, 1972). In this case, the "interpreter" is the set of metal fingers pressing against the player roll that sense the holes and activate, by air pressure, their corresponding keys. Clearly, the interpreter can be quite "mechanical" and is not an occult entity. Whether or not one imputes physiological "reality" to his program, our external knowledge of the program gives us predictive power regarding the behavior of the device it simulates or controls.

2. The Executive Monitor

In a hierarchical program, the top-level routine is called the *executive,* and this is a useful concept in models of the cognitive system. The executive calls routines at the next lower level and keeps track of where these subroutines are to return their results. The executive monitors the number of subgoals being generated using a particular method (say, in a problem-solving situation). It also evaluates (from feedback) how the current plan is progressing. The executive may interrupt and switch to another subgoal either if that other one suddenly becomes more important or if the current method of attack on a subgoal seems not to be progressing satisfactory, for example, if it exceeds a work limit. The executive also notices when a subgoal has been completed so that its results may be used in selecting the next step or next goal to be worked on. The executive is itself just another bit of program and is quite deterministic in its actions.

3. Plans for Learning

The primary concern of this chapter is with learning and remembering, so we shall usually be discussing the operations of the cognitive system when the current plan controlling activity is either a plan for learning (acquisition of specified information) or a plan for remembering (retrieving previously stored information). In characterizing learning, one is interested in analyzing the means by which the perceptual process enters information into short-term memory, and in how that information comes to be represented and related to other things as it is stored in long-term memory. For example, we may regard memorizing as a problem-solving task for which the person selects and applies particular "mnemonic strategies" or plans which have proven useful to him for solving similar memorization problems. These plans for deliberate memorization (for example, concentration and verbal rehearsal) are clearly learned and develop along side other intellectual abilities of the child (see Meacham, 1972). It is clear

too that the degree of learning we observe depends heavily upon which mnemonic strategy the person employs when he is exposed to the material.

With respect to plans for remembering, the significant issues to be addressed are (1) how the current "test question" or problematic situation can retrieve from memory information relevant to it, (2) what is the organization or "filing index" of the storage system, (3) what search procedures are carried out to locate relevant memories, and (4) how do we construct answers based on the fragmentary information we retrieve.

II.   INFORMATION-PROCESSING ANALYSIS

A.   Stages and Components

The information-processing approach assumes that perception and learning can be analyzed conceptually into a series of stages during which particular components ("mechanisms") perform certain transformations or recodings of the information coming into them. The subject's eventual response (for example, the perceptual judgment "I see a giraffe") is considered to be the outcome of this lengthy series of operations. Each stage in the system receives as input the information as coded in its predecessor stage, operates upon it so as to condense, abstract, recode, or elaborate it, and then passes this product on to the next stage in the analysis. Since external stimuli can not get inside an organism, the representation of them ("internal symbols") and their interrelations ("symbol structures") is what we call "information," and this is the content we describe in our theories. As experimenters, we try to devise techniques for studying the representation of information at each stage, the nature of the recoding done at that stage, the experimental variables that influence the duration of the stage, and so on. We try to represent the supposed causal order of these processes in terms of a flow diagram in which blocks represent component processes occurring at successive intervals. Each block is labeled according to its typical function. In terms of the computer analogy, the operation of each processing stage would be represented by a subroutine. The identification of these processes and their causal operation remains the major, unfinished task for information-processing theories.

B.   Some Concepts of Information Processing

However, before launching into a description of the "human system," it is prudent to mention briefly some of the terminology to be used as we proceed. According to Newell and Simon (1972, p. 20 ff.), an information processing system (IPS) consists of a sensory system, a response generator,

a memory, and a central processor. They then provide the following set of interrelated definitions:

1. There is a set of elements, called *symbols.*
2. A symbol structure consists of a set of *tokens* (equivalently, *instances* or *occurrences*) of symbols connected by a set of *relations.*
3. A *memory* which is a component of an IPS capable of storing and retaining symbol structures.
4. An *information process* is a process that has symbol structures for (some of) its inputs or outputs.
5. A *processor* is a component of an IPS consisting of:
    (a) a (fixed) set of *elementary information processes* (eip's);
    (b) a short-term memory (STM) that holds the input and output symbol structures of the eip's;
    (c) an *interpreter* that determines the sequence of eip's to be executed by the IPS as a function of the symbol structures in STM.
6. A symbol structure *designates* (equivalently, *references* or *points to*) an object if there exist information processes that admit the symbol structure as input and either:
    (a) affect the object, or
    (b) produce, as output, symbol structures that depend on the object.
7. A symbol structure is a *program* if (a) the object it designates is an information process, and (b) the interpreter, if given the program, can execute the designated process. (Literally this should read, "if given an input that designates the program.")
8. A symbol is *primitive* if its designation (or its creation) is fixed by the elementary information processes or by the external environment of the IPS [Newell & Simon, 1972, p. 20].

Also,

Reading or encoding consists in creating in memory internal symbol structures that designate external stimuli; *writing* is the inverse operation of creating responses in the external environment that are designated by internal symbol structures [Newell & Simon, 1972, p. 21].

An illustrative example from Newell and Simon (1972, pp. 21–22) is a program for translating Morse code. Thus, an external stimulus such as $(-, -, -)$ would be read (recognized) in short-term memory as a symbol structure, namely, three tokens of the dash symbol connected by the "next" relation. (The "next" relation is used in building up lists of symbols.) Next, the designation of this symbol structure would be "looked up" in long-term memory. For this example, the corresponding symbol structure in long-term memory is attached to the letter S. Hence, the $(-, -, -)$ structure in short-term memory would be replaced by its designation code, S. As the receiver listens, he builds up an *encoded* symbol structure in short-term memory such as (S, I, T), which itself in a later pass designates a word. This discussion describes stimulus recognition at intake.

The process of sending or writing (generation) would just use the machinery in the inverse order.

C. Elementary Information Processes

There would be a set of elementary information processes that can be executed by the system. Among them would be one to "find the member on a list that is next to a specified one"; another one would "check whether two symbol tokens, x and y, are identical." By combining these two elementary processes, one can synthesize a larger process, such as one which compares two lists of symbols to see whether they are identical.

A "symbol" in such a system is a unit which permits access to any information associated with it: this information might be its "referent" or another symbol structure. For example, the internal symbol "John Smith" might refer to—be associated with—a symbol structure describing his appearance and other facts about him.

A major claim of computer science is that there is a relatively small set of elementary information processes which will suffice to produce the full generality of any symbol-manipulating and problem-solving activity we could want. Newell and Simon suggest that the following ones comprise a sufficient set:

1. *Discrimination.* It must be possible for the system to behave in alternative ways depending on what symbol structures are in its STM. Furthermore, the behavior needs to be *arbitrarily* alterable; i.e., transfer of control to an independent program must be possible.
2. *Tests and comparisons.* It must be possible to determine that two symbol tokens do or do not belong to the same symbol type. Often comparisons are directly coupled with conditional behavior, but they may equally well lead to the production of a conventional symbol (e.g., *true* or *false*) that can later be discriminated.
3. *Symbol creation.* It must be possible to create new symbols and set them to designate specified symbol structures. Again, this process must be performable arbitrarily; i.e., whenever a new symbol is desired it can be created, and it carries no meaning other than that it designates the desired symbol structure. Whether the system must also be able to destroy symbols depends primarily on whether memory capacity is limited.
4. *Writing symbol structures.* It must be possible to create a new symbol structure, copy an existing symbol structure, and modify an existing symbol structure, either by changing or deleting symbol tokens belonging to the structure or by appending new tokens with specified relations to the structure. Many variations are possible, as long as they permit building up arbitrary structures.
5. *Reading and writing externally.* It must be possible to designate stimuli received from the external environment by means of internal symbols or symbol structures, and to produce external responses as a function of internal symbol structures that designate these responses.
6. *Designating symbol structures.* It must be possible to designate various parts of any given symbol structure, and to obtain designations of other

parts, as a function of given parts and relations. Again, this can be achieved in many ways but it must be always possible; i.e., there must not be any parts of symbol structures that are in principle inaccessible.

7. *Storing symbol structures.* It must be possible to remember a symbol structure for later use, by storing it in the memory and retrieving it at any arbitrary time via a symbol structure that designates it. How much memory is available, of course, conditions strongly how complex the totality of stored structures may be. The memory must be highly reliable over time [Newell & Simon, 1972, pp. 29–30].

It is supposed that such components as these will be needed in building a model to simulate some interesting behavior. Assuming the existence of such information processes does not commit one to any particular level of analyzing a given psychological phenomenon. What may be primitive symbols at one level of analysis (for example, words for language understanding programs) are large symbol structures in models aimed at a more molecular level of analysis (for example, words as ordered sets of graphemic features in models of word perception). The content of the above processes will also vary with particular applications—for example, the "external responses" that can be produced (under 5 above) obviously differ depending on whether one is modeling a crayfish, an infant, or a human adult.

Whatever the phenomenon being modeled, it is assumed that most of the above processes will come into play at some level of analysis. To illustrate, in modeling the memory-scanning task of Sternberg (1969), one might assume that presentation of the stimuli "7, 9, 3, 6" as the memory set causes the subject to read these in (Process 5 above), to match them to symbol types in long-term memory (Processes 1 and 2), and set up copies (Process 3) of these symbol types in a symbol structure in short-term memory (Process 7). Thus, a element called "Most recent list" is made to designate the symbol structure "7–next–9–next 3–next–6" (by Processes 3 and 4). The test probe, "Is 3 in the most recent list?," is read in (Process 5), compared symbol by symbol to successive elements on the most recent list (using Processes 1, 2, 6), with a "yes" response being generated (Process 5) because the test probe matches one of the items in the set. The operation of "get the next symbol on the list" is an elementary process (6 above) as is the process of "compare list symbol to probe symbol for a match" (2 above). We thus see how a given model analysis uses such information components as these.

III. OVERVIEW OF THE COGNITIVE SYSTEM

With this introduction to basic vocabulary and concepts of the information processing approach, let us turn to sketching the viewpoint as it applies to the behaving organism. As noted above, the approach tries to represent
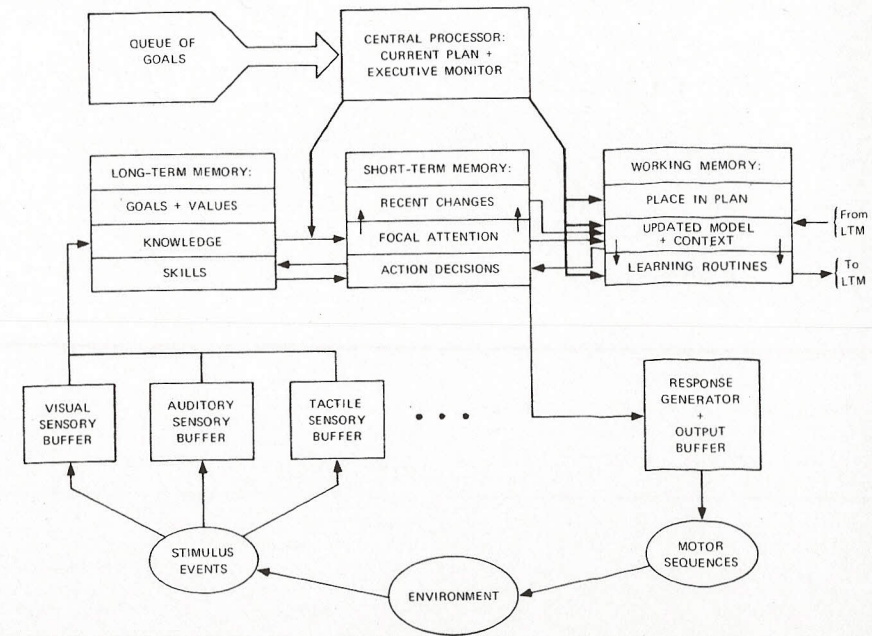
FIGURE 2   Diagram of the principal components of the perception and memory system. Arrows indicate the direction of information flow and control among the components.

the cognitive system in terms of a flow diagram in which blocks represent components and successive stages of processing. Although many such flow charts have been proposed, a sort of "modal" summary chart might look like that in Figure 2. The interconnections between components are specified by the arrows that suggest the direction of action and control. The components may be divided roughly into the sensory system, the response system, the long-term memory system, and the central processor wherein occurs the active processes coordinate with memorizing, thinking, evaluating, and decision making. We now describe these components in Figure 2, starting at the sensory end.

A.   Perception

1.   Sensory Buffers

The environment provides an array of patterned stimulus energies, reflecting both static settings as well as dynamic changes ("events"). Our sensory receptors have transducers which convert these analog energy patterns into a digital code (that is, nerve firing patterns). Apparently at a very low level in the sensory system (for vision, the retinal ganglion cells), certain elementary "preattentive" processes come into play: in vision

(which will always serve as our illustration), these processes perform contour enhancement, centering of the stimulus, segregation of the figure from the background, and so on. Quite a bit is known about the neural mechanisms involved at this stage (see Ratliff, 1965; von Békésy, 1967).

As the stimulus pattern is registered on the receptor surface (say, the letter H projected on the retina), a set of feature detectors located along the visual pathways begin extracting significant features from the stimulus. Features of H would be that it is black, has two vertical lines, one horizontal line, is open at the top and bottom, etc. These features are apparently extracted in parallel for all patterns falling near the fovea. The feature coding at this level is quite automatic and seems little affected by momentary expectations. The physiology underlying and corresponding to feature detection has been intensively investigated, steming from the classic work of Hubel and Wiesel (1962, 1968) and many others. We now know how to construct plausible neuronal-network models that will extract a large and relatively sophisticated set of features such as the size, orientation, direction of motion, and velocity of movement of lines, edges, arcs, and so on (see, e.g., Lindsay & Norman, 1972, Chapter 2). It appears that the connecting up of the neural elements which extract a given feature (for example, horizontal lines) requires that the developing infant receive normal visual stimulation, since restricting vision during rearing (say, to only vertical lines) causes selective deficits in appropriate cortical detectors in kittens (see review by Blakemore, 1974). Restrictive environments aside, Eleanor Gibson (1969) has argued that perceptual development should be viewed in terms of the organisms's growing sensitivity to more, and more subtle, distinctive features of his perceptual environment.

If the stimulus is a brief change (say, a letter flashed in a tachistoscope), the stimulus is held briefly in a sensory memory as an image or icon while its features are being extracted. We represent this by saying that the image is held in a "visual buffer" as a sort of fleeting memory of the physical stimulus. The duration of the icon can be varied by the field-to-flash contrast before and after the flash. The icon can be "washed out" by a bright masking flash just after the shape is flashed, or it can be replaced by a second form being flashed very soon afterwards to the same retinal location (or to the homologous location on the other retina, see Turvey, 1973). The brief duration of the icon severely limits the number of features which can be extracted from a brief glimpse. In typical conditions the icon lasts about a quarter of a second which, interestingly, is also about the time for a saccadic eye movement. This suggests that as the eye sweeps over a scene, the icon from view $n + 1$ replaces (rather than becoming confused with) the icon from view n. Of course, as we scan several times over a scene, we pick up progressively more information, which is integrated into a more complete representation of the scene in consciousness.

It appears that such sensory analyzers and buffers exist for all the major sense modalities—for vision, touch, audition, and so on. Just as we have fleeting visual icons, so do we also have fleeting acoustic and tactile images. These sensory buffers operate in parallel with little interference between them. The interference that arises between monitoring two modalities appears to occur "higher up," in the conscious processor.

2. Pattern Recognition

The features extracted from a given stimulus object comprise a sort of coded "description" of the object. In computer jargon, it is a list of attribute–value pairs. In terms of the theory proposed by Anderson and Bower (1973, nicknamed HAM, an acronym for Human Associative Memory), a feature list is a *conjunction of propositions* about the object. For example, the letter E might become coded as "this object consists of 3 horizontal lines which make left-contact with a vertical line." It will be noted here that "feature" has now been extended in meaning to include not only presence or absence of a line but also relationships between the lines and angles and their type of connectivity. Again, work on automatic character recognition (see Duda & Hart, 1973) has provided a variety of mechanical algorithms (in some cases, neural networks) for recognizing such simple relationships among parts.

Following this feature extraction stage, there is an *identification* stage during which the system tries to decide how to classify the stimulus object. In brief, this is assumed to occur by a weighted matching of the current feature list against a likely set of prototypes in long-term memory, with the input being classified according to the name of the best matching prototype (see Reed, 1972). Identification accuracy depends upon the quality or extent of the sensory information extracted in the first stage. Performance also varies with the number of alternative classifications that are expected or are likely given the context. This has been demonstrated many times in degraded perceptual situations in which the size of the expected identification set is varied: accuracy is higher, the fewer the alternative classifications the subject is deciding among (see Garner, 1962). A meaningful stimulus will also be identified more easily if an associated semantic context has been activated near the time the degraded stimulus arrives. Thus, I can more accurately identify a picture of a loaf of bread in a brief glimpse if I am primed to think of kitchen-related items. Or I am more likely to identify the word "America" if I know beforehand that the item fits into the phrase "the United States of _____." Identification accuracy increases directly with the likelihood that one could guess the target item on the basis of the context alone.

Morton (1969) presented a model suggesting how sensory and contextual information might combine in such situations. He supposed that familiar patterns (such as words) have corresponding central units (called *logogens*) accessed directly by the sensory feature lists. Thus, there would be a logogen for common words like EGGS, KEGS, CAKES. By some unknown means, the prevailing semantic context can also supply excitation to particular logogen units. If I have just heard "For breakfast, I had ham and _____" just before I hear a very noisy signal (of "eggs"), the latter signal acting alone might activate logogens that are acoustically similar to EGGS, such as KEGS, LEGS, ACHES. But the context has already placed excitation on a number of logogens related to breakfast foods. The logogens sum their excitation from the two sources, and the dominant one makes its response available to consciousness. Thus, the person will hear the signal as EGGS. The logogens in Morton's theory could be identified with "word nodes" in HAM, the Anderson and Bower theory; that is, a word like EGG would have a central node in the memory network corresponding to a proposition describing how the word is spelled, pronounced, and what it sounds like. What is not specified in either Morton's or HAM's account is the "comprehension machine" that can use a semantic context to activate selected words.

3. Segmentation and Scene Analysis

The character identification process described above applies to a single visual object like a printed letter. Having taught the machine to identify single letters, we do not want it to boggle when we show it strings of letters. We want it to treat this pattern as an ensemble of known subparts rather than as an entirely new object. To do so, the pattern recognizer must have the ability to segment or decompose a scene into its several primitive objects, and arrive at a proper identification of these and their interrelationships. Though illustrated for visual scenes, the segmentation problem is just as troublesome for speech recognition since speech tends to flow together in a stream of overlapping phonemes.

The segmentation problem has not been solved in general, but there are several significant programs in artificial intelligence work which appear to be very promising models. To achieve segmentation, a series of passes and interactions are clearly needed between different levels of analysis—with lower-level visual processes that are building up information about elementary lines, angles, joints, etc., and higher-level processes that are testing the developing lower-level descriptions for coherence and plausability. The SEE program of Guzmán (1969) was one of the earlier sophisticated scene analysis programs. It analyzed pictures of three-dimensional blocks scattered about in various orientations and occlusion relations.
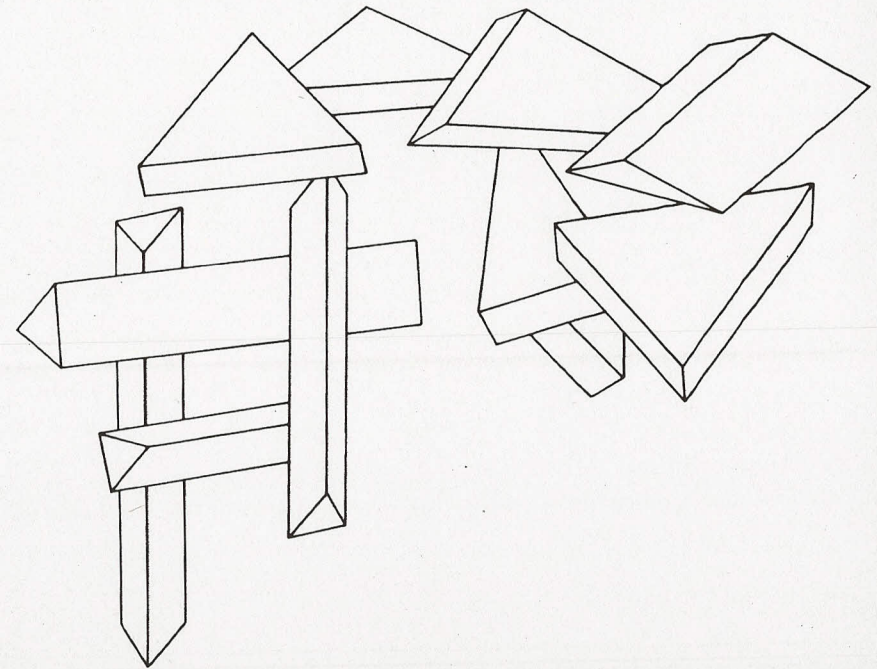

FIGURE 3   A typical toy-block scene analyzed by Guzmán's SEE program. (Reproduced by permission from Guzmán, 1969.)

Figure 3 shows an example scene analyzed by the SEE program. From analyzing a number of such scenes, Guzmán concluded that the most important information arises from intersections, where the surface of one object changes or where the surface of one object is occluded by another surface. Figure 4 shows several of the significant local signs identified by Guzmán, each of which supplies particular types of clues for interpreting the scene. For example, the L intersection suggests that the surface to the left (1 in Figure 4a) belongs to a different body from the surface on the right (2). An *arrow* usually suggests two bodies, one with the two surfaces 3 and 4 and the other with surface 5. A pair of complementary T joints strongly suggests occlusion of one object (with surfaces 7 and 8) by a foregrounded object (with surface 9). Such local information gathered from vertex clues is combined with more global rules and assumptions about visual scenes in order to arrive at a meaningful segmentation and identification of the parts of the scene.

Winston (1970) has gone beyond Guzmán's program in developing a system that not only segments a visual scene but also describes its objects in terms of geometrical concepts and their interrelations. Thus, Winston's program deals with architectural notions like those in Figure 5. Each of
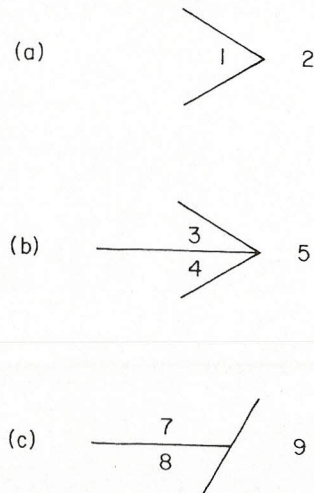
FIGURE 4 Examples of local signs which, in local patterns, signify surfaces, angles, corners, objects, and occlusion relations within block scenes. (a) An L-vertex where two lines meet; (b) an arrow, where three lines meet at a point, with one of the angles bigger than 180°; (c) a T, of three concurrent lines with two of them collinear. (Adapted from Guzmán, 1969.)
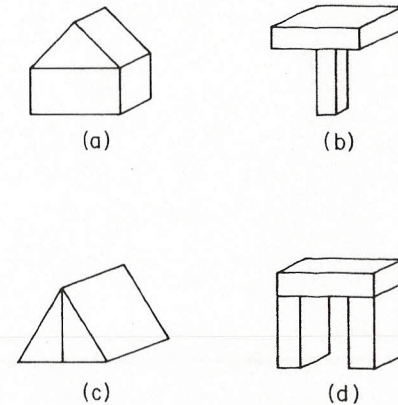

FIGURE 5 Example geometric structures or concepts which Winston's program can recognize and deal with as units: (a) a house, (b) pedestal, (c) tent, (d) arch.

these concepts has an elementary structural description. For example, an ARCH is characterized as having three blocks, two upright blocks that are not abutting and are supporting the third. Such structural descriptions (which are labeled relations among elementary concepts) are the output of the perceptual analysis performed by Winston's program. Such descriptions also illustrate the way in which these concepts (of an *arch, tent, house,* etc.) are learned and stored in memory. A test object is classified by matching its structural description to the set of concept descriptors in memory, selecting its name according to the best match (if any).

An important feature of such intermediate-level concepts is that they simplify enormously the description of a scene. For example, consider the scene composed of a tent to the left of a house, with both sitting on top of an arch (see Figure 5). There would be at least seven simple objects—bricks and wedges—in that overall scene, and they would exhibit very many interobject relationships. *Before* the system learns the intermediate-level concepts in Figure 5, all of these parts and myriads of relationships would have to be entered into the description of the scene. However, *after* having learned these concepts, they can be used as units to simplify the overall description to something approximating our English proposition, "a tent to the left of a house both supported by an arch." We can thus see the tremendous economy of internal description provided by our learning of higher-order concepts.

The Winston program is just one of several scene analysis programs available; others exist in conjunction with various "hand–eye" projects and robot projects at several American research centers. Scene analysis is a very difficult problem, and at present only relatively simple scenes are being analyzed. But it is hoped that the methods found successful there will

generalize to more complex scenes. Since there is no psychological theory of how people perceive full scenes, these programs are the "best" theories available by virtue of being the *only* theories close to the computing power needed to do the job. Similar progress is being made on speech perception where the segmentation problem is difficult but is being attacked by a "top-down" analysis in terms of semantic constraints (produced by the topic of conversation) on what words are probable in the current context (see Newell *et al.,* 1973; Vicens, 1969).

B. Short-Term Memory

1. General Characteristics

In the course of identifying a stimulus pattern, an internal symbol for it is brought into an active state. Alternatively, we may say that the corresponding symbol has been entered into the subject's short-term memory. Short-term memory (STM) is the active part of the central processor that holds the symbols currently in the focus of attention and conscious processing. Short-term memory need not be viewed as a "place," "register," or "box" physically distinct from long-term memory; STM and long-term memory may only be two different states (or operational characteristics) of the same memory network. For example, the items in the "STM state" may only be that subset of internal symbols (or memory nodes) that have been recently activated. In a computer simulation program, however, entering an item into short-term memory would correspond to placing a token (copy) of its recognized symbol type onto a short list named STM; an alternative representation would be to enter a temporary association from a node named STM to the address of the just recognized symbol type.

The basic characteristics of STM are as follows: (1) as noted, it is the *active* partition of memory; (2) the processor has fast access to the items in STM (that is, retrieval of items from LTM is slower); (3) STM tends

to preserve the temporal order of incoming symbols; (4) STM has a severely limited capacity, about 4–6 symbols or coded items.

Let us comment briefly on these latter three properties. The point about access speed is revealed in experiments showing that someone can respond more quickly to items or facts he has just been thinking about. A likely explanation is this: when a given item has been thought of, its "neural" logogen continues for a while in a state of excitation, with the result that a subsequent input of the item will cause the logogen to pass the excitation threshold needed to fire more rapidly, thus making the item rapidly available to consciousness. This is a plausible mechanism for various sorts of "priming," wherein activation of an item brings it temporarily into readiness to fire again (or "lowers its threshold"). The assumption about access speed in STM is brought out in studies by Sternberg (1969) and others (reviewed elsewhere in these volumes) investigating the time it takes a subject to decide whether a probe or test item belongs to a small predefined set of symbols held in STM. The times increase with the number of symbols in the set, suggesting some kind of a serial scan by the processor over items in STM. While this indicates that the processor does not have immediate access to all items in STM, it is still the case that items in STM are responded to more rapidly than those in long-term memory (e.g., Wescourt & Atkinson, 1974).

That STM preserves the temporal order of arrival of items is clear whenever one is asked to "repeat back" in any order a short string of verbal items he hears or sees. The overwhelming tendency is to preserve temporal order. If one asks for a word series like "eagle, chair, sparrow, soft" to be reported back in an order determined by a semantic classification (for example, furniture, then birds), or another series is to be reported according to rhyming relations, the output time is much slower. It is as though the person has to recirculate, classify, and reshuffle the items slowly to do these latter tasks.

That STM has limited capacity is obvious in terms of the small number of "things" we seem able to keep in mind all at once. Also, our ability to repeat back a random string of items (the "memory span") is severely limited to about seven items. In theory, the span should be longer than the STM capacity since some of that recall is aided by long-term memory (about the test series) build up during initial intake of the series.

2. Chunking

The capacity of STM is best stipulated in terms of items or "chunks" of information. G. Miller (1956) first called attention to the fact that verbal memory span is relatively constant, regardless of the "size" of the vocabulary from which the symbol units are being drawn (for example,

letters, digits, syllables, colors, words). But what is a chunk or "unit" for the system? The chunk can only be defined circularly as a stimulus pattern or sequence which the perceptual system recognizes as a familiar, single unit for which there is an internal code (or "symbol"), possibly one designating further semantic structure. Thus, PEN is treated as a unit rather than as three letters, whereas ENP is just three letters. Since JACKPOT has units corresponding to the seven single letters, to four spelling patterns, to two single words, and to one compound word, what determines the "level" of the units that will be entered into STM in this case? Under standard operation, the system normally takes the simplest description of the object (that is, fewest symbols); in this case, that is the internal symbol corresponding to the compound word—the highest level or largest stretch of input over which a successful match occurs. By instructions to the subject, of course, we could bias his "encoding" so that he only pays attention to the letters or the lower-level units.

The units that are identified by the perceptual system depend partly upon the physical characteristics or "gestalt groupings" of the entering elements. The sensory system groups together elements that occur close together in time (or space) and which have similar appearances. Thus, a temporally ordered series of letters with pauses like IBM–FBI–CIO–TV will be remembered much easier than the series IB–MFB–ICI–OTV. Pauses segment the series into four familiar acronyms in the first case but not in the second. But why, really, it this important for memory? It creates a difference in the complexity (or numbers of linkages) of the symbol structures which must be set up in STM to encode and remember the different strings. The well-segmented string eventually is encoded as a string of four tokens of known patterns, whereas the poorly segmented string is represented as twelve tokens (for the individual letters). Therefore, the former string is remembered more easily because it has many fewer connections to be learned.

3. Coding of Information in Short-Term Memory

One of the standard strategies that people have evolved to deal with memorization tasks is to *recode* or translate the recognized symbols in STM from one code domain to a second code domain, hoping thereby to improve their memory for the material. The general model of this coding and decoding is diagrammed in Figure 6. For example, the symbols (or symbol structures) in Domain 1 might be Spanish words that are to be coded into German, or they may be visual letters to be coded into manual sign language, or Morse Code into letters, or printed sheet music into performed music. A common practice seen in verbal learning is for the subject to convert visually presented items into their names. A well-known consequence of this translation is that items then tend to be confused and to
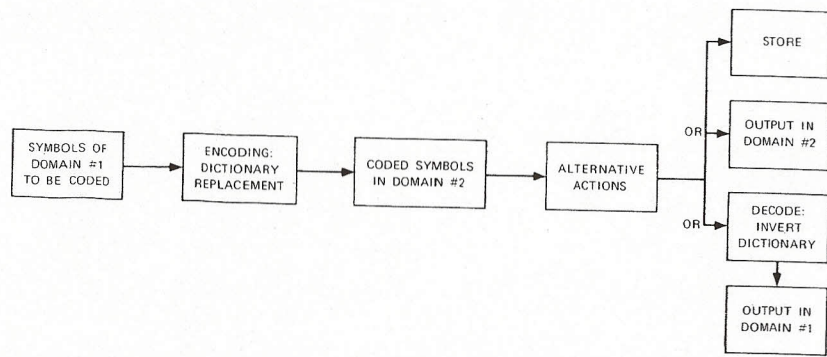
FIGURE 6   Schematic diagram of the encoding process, with the coded symbols then leading to alternative actions depending upon the task demands.

interfere with one another according to their phonetic similarity rather than their visual similarity (Conrad, 1964). If verbalization of the presented items is prevented, then the person seems to rely much more on the visual appearance of the item (see Estes, 1973; Parkinson, Parks, & Kroll, 1971). This suggests that STM can hold symbols (for the order of seconds) that are distinguished by their modality of presentation (for example, auditory or visual).

Different coding schemes may be classified according to whether they preserve, add, or subtract symbols. Some of the mappings are one to one, as indicated above: that is, one visual word converts to one phonetic word, one picture to one name, one sound to one label, and so on. Other mappings are reductive, either because they translate a long string of symbols of Domain 1 into a single symbol in Domain 2 (for example, octal coding, Morse coding, letters to word) or because they select only one distinguishing fragment from the input and discard the remainder. An example of the former kind is when binary sequences like 11001100 are converted into rule-like descriptions of the form "alternating pairs of 1's and 0's (see Glanzer & Clark, 1963). Examples of fragmentary coding would be (1) labeling pictures of different breeds of dogs by the superordinate category, thus discarding distinguishing features, or (2) classifying words according to whether or not they have an initial "bo-" sound, thus aggregating together bow, bowl, boat, etc. Memory for an input may be enhanced by coding to the extent that the code is "shorter" than the original, and permits accurate reconstruction of it. The Glanzer and Clark experiment is a case in point: their subjects were better at recalling those binary series that evoked shorter descriptive codes. In case the code is obtained by ignoring some of the input, then memory for that input will be confused and indeterminate, specifically regarding variations in the ignored parts of the original pattern.

The third class of mappings is elaborative, in which the number of elements appears to be increased during the translation. An example would be coding the name of an event (for example, VE Day) in terms of a set of memories one has associated with that name; another is replacing a word by an image of a known referent and the context of one's knowledge about that (for example, recoding "barber" by imaging the last barber you visited). An equivocal case is so-called "natural language mediators" (NLMs) that adults concoct to remember meaningless materials like nonsense syllables (see Prytulak, 1971). In order to remember, subjects will add or change letters to convert the syllable into a meaningful word. Thus, LOS converts to LOSe, LOM to LOaM, TAY to sTAY, and so on. Prytulak argued that adults use this coding as a memory aid, since a three-letter symbol structure (L–next–0–next–M) is thus reduced to a meaningful and shorter two-symbol structure (LOAM–without–A). Prytulak supplied a variety of evidence in support of this memory coding, especially for syllables presented at a slow rate.

These several coding schemes generally take time—initial symbol identification, looking up its code in the dictionary, substitution of the code for the original symbol. Consequently, such coding typically occurs only when time is available. But such coding is a trainable skill that improves with practice. Reading—converting the printed to the spoken word or to a semantic meaning—is surely one of the most advanced coding skills we have.

4. Rehearsal Processes in Short-Term Memory

Because most investigations of STM are conducted with adults being exposed to verbal items, there has been much concern in the experimental and theoretical literature with verbal rehearsal (overt or covert) as a subject-controlled strategy for memory enhancement. It is widely agreed that some mechanism akin to subvocal naming of a small series (1) can be consciously initiated, (2) can selectively edit or delete certain elements from study, (3) can maintain a small set of verbal items in STM, and (4) under appropriate conditions can cause information about the rehearsed items to be laid down more permanently into long-term memory.

These assumptions have been summarized most explicitly in the theory of Atkinson and Shiffrin (1968, 1971) and Shiffrin and Atkinson (1969) though anticipated by several others (e.g., Bower, 1964; Broadbent, 1958; Waugh & Norman, 1965). Short-term memory is identified with a rehearsal buffer which is said to have a fixed capacity of r symbols. In a memory task in which the subject is to study and learn a list of unrelated items presented rapidly one at a time (for example, a word list for free recall, or a continuous paired-associate list), the theory supposes that the person

takes successive items into his rehearsal buffer and rehearses them in order to learn them. Each rehearsal cycle is assumed to transfer to LTM some information about the items in STM. Once the arriving stream of items fills up the r slots in the buffer, a new arrival causes a prior member of the buffer to be "bumped out" and replaced by the new arrival. At that moment, the system will stop learning anything more about the bumped-out item; its sole remnant in the system is whatever information may have been built up about it in long-term memory during its time of residence in the rehearsal buffer.

This leads to an "occupancy model" of STM; STM is analogous to a luncheon counter with only r (say, four) seats where customers (items) arrive for servicing, and a set of clocks in long-term memory record how long each customer is allowed to sit before he is bumped off his stool by a new arrival. In a recall task (say, free recall of the list items), it is presumed that an item can be recalled with probability one if it is still in the rehearsal buffer; if it is not in STM, then it may be recalled with a lower probability that is greater the longer its residence time in the buffer (or the greater the amount of information stored in long-term memory). The buffer theorists have never been very explicit about exactly *what* was being stored in long-term memory (they give a functional definition—"it is whatever enables the person to recall better"). In the propositional theory proposed by Anderson and Bower (1973), what is being stored and growing in strength is a proposition of the form, "symbol X appeared with symbols Y and Z in the context of experimental list N." It is then a cluster of such propositions that are accessed when the subject is later cued with the command, "recall all symbols that occurred in the context of list N."

5. Differential Rehearsal and Priority Setting

However the model is formulated in detail, it is clear that variations in rehearsal usually lead to differences in learning and performance. It is assumed that selection of a rehearsal strategy and of the set of items to be rehearsed is a plan which is under control of the executive which is trying to maximize achievement of particular goals. Let us briefly mention a few learning effects which apparently can be explained by differential rehearsal.

a. *Intentional versus incidental learning.* In most measures of recall, intentional learners typically exceed incidental learners. Recording of what they do during exposure to the materials shows that intentional learners are actively rehearsing more. Therefore, motivation to learn and incentives for learning affect what the subject does with the material at the time of exposure, and his later performance is largely determined by those activities.

b. *Selection of a functional cue.* In learning with a compound stimulus (for example, a picture of an object and its name), the subject will select one or another element to rehearse depending on how he expects to be tested. If we show him a detailed drawing of a cup, he will look at and rehearse different things if he expects to have to free recall the names of all the objects in the study set than if he expects to have to select exactly that cup from among a set of similar cups (see B. Tversky, 1973).

c. *Rewards and punishers as information.* We can vary the priority of rehearsal that a subject assigns to two items by telling him at the time of presentation that recalling one will be worth a lot (say, a dollar) whereas recalling the other will be worth very little (a penny). The central processor gives priority to the high-payoff item; it is kept in the rehearsal buffer and newly arriving material tends mainly to displace (and cause forgetting of) the low-payoff items. This analysis is bolstered by the fact that people can remember the approximate payoff they have been promised for remembering a given item. It is likely that the effect of rewards and punishers given after the response in a Thorndikean situation (see the chapter by Nelson, next volume) is largely due to the information they provide about what is to be rehearsed and what can be ignored (see Estes, 1969).

d. *Intentional forgetting.* Many experiments (reviewed by Bjork, 1972, and Epstein, 1972) have shown that subjects can selectively "delete" items from their current STM. This can be done simply by telling the learning subject soon after he has entered one or more items into STM that he will no longer be held responsible for remembering those items (that is, he can "forget" them). This forget instruction has several effects: it causes the subject to cease rehearsing the forget items and to devote more rehearsal to the items on which he expects to be tested; the forget signal is also a stimulus event that can become associated with any items with which it was paired, so that the person may be able later to indicate whether a given item was a "forget" or a "remember" item; and earlier "forget" items may be rejected as possible associative contacts to help organize later "remember" items in a free recall list. The literature cited may be consulted to become acquainted with the experimental data surrounding investigations of intentional forgetting.

e. *Priority of distinctive items.* A result by von Restorff (1933) and many others is that an item that is distinctively different or surprising in relation to the other items in a learning series will tend to be learned first, but it does so at the expense of neighboring items of its series. An experiment by Ellis, Detterman, Runcie, McCarver, and Craig (1971) illustrates

the point: subjects saw many series of slides of simple objects, recalling their names after each series. For the critical series a shot of a pornographic nude was shown without comment in the middle of the picture series. Plotting serial position curves for free recall, the nude picture was recalled 100% (well above the corresponding picture in the control list), whereas pictures adjacent to the nude picture were recalled worse than their corresponding controls. An explanation in terms of the buffer model is that the startling item caused rehearsal of just-prior items to cease, whereas subsequent items were either not entered into STM or were assigned low-priority rehearsal scores because the "nude picture" code remained in a state of high activation for awhile.

f. Instructional differences in rehearsal. The experimenter can instruct subjects to rehearse the series of study items according to distinctively different patterns, for example, one at a time, by nonoverlapping triplets, by cumulative addition of the next (see Atkinson & Shiffrin, 1971). These instructed variations in rehearsal produce differences in later free recall as predicted by the buffer model (for example, reduced primacy effect in "one at a time" rehearsal, enhanced primacy in cumulative rehearsal, and triplet-clustering in the third condition).

g. Individual differences in rehearsal. There are large differences among individuals in memory. Young children, by and large, remember less well than older ones; children with high IQs usually remember more than those with low IQs—which is why most IQ tests include memory subtests. There are doubtless a variety of causes for these differences. One difference, for example, is that the older, more educated person may have become familiar with more subsequences (chunks) of the test materials, so that he can recode some lengthy series into a slightly shorter sequence of internal symbols (for example, letters to spelling patterns). Another likely contributor to the difference in memory is the fact that older, smarter children use more efficient mnemonic methods. Flavell and his co-workers (e.g., Flavell, Friedrichs, & Hoyt, 1970) have found marked differences among children in their spontaneous use of verbal naming and sequential rehearsal of object series they were to remember; older children tend to rehearse more and remember more. If younger children are induced to rehearse more efficiently as they study the series, then their later recall equals that of the older children. In other work, it has been found that nursery school children do not behave differently under instructions to just "look at" versus "memorize" a picture series, whereas older children employ deliberate memorizing strategies when given the latter instructions. This suggests that memorization strategies develop gradually as a way to deal with a future contingency, one that small children do not yet understand (see Appel, Cooper, McCarrell, Sims-Knight, Yussen, & Flavell,

1972). Similarly, Belmont and Butterfield (1971) found marked differences in the way normal children and retarded children pace themselves through a self-presented list of items they are to memorize. Normals often do "grouping" and "cumulative rehearsal," returning repeatedly to the beginning to go over the entire series up to the present point as though testing themselves. In contrast, retarded children tend to zip rapidly through the study series without substantial pauses. There are correspondingly large differences in the two groups' recall of the series, particularly in the early ("primacy") items of the series. If the retarded children are trained to use the cumulative rehearsal strategy, with longer pauses for self-testing, their performance improves to become much more like that of normals on this task. This research area, comparing children of different mental ages on their mnemonic methods has just begun to be explored in depth over the past few years. The practical hope, of course, is to find teachable methods that will accelerate memory development in small children and retarded children.

6. Shortcomings of Verbal Rehearsal Models

The foregoing is just a partial list of the sorts of results that are explicable by the notion of a verbal rehearsal buffer that is under strategic control. For other examples, see Atkinson and Shiffrin (1968) and especially Rundus (1971), who devised the method of making rehearsal overt in order to examine whether the mnemonic influence of several variables (such as the spacing between presentations of a repeated item) was mediated through their direct effect on amount of rehearsal of the item in question. These methods have indeed supplied us with very valuable information.

However, there are several shortcomings of the supposition that short-term memory is coextensive with the verbal rehearsal buffer as described in these models. A first difficulty arises when we realize that we have millions of memories for nonverbal events or items which we have not verbally described—for paintings, music, scenery, all varieties of sensory magnitudes such as the pitch and loudness of tones, the length and grayness of lines, pressures, smells, tastes, etc. Of course, some of these we do classify verbally (for example, "it was about a three-inch line"); such a verbal description clearly can be rehearsed, and it will influence our memory (or reproduction) of the stimulus. But it is just as obvious that we can remember something about nonverbal events even if we do not (or can not) describe them. They seem to be represented in terms of analogical structures that are not necessarily connected to the verbal production system.

It is a moot question which modalities have a usable covert rehearsal mechanism. It would appear that some kind of motor reproduction system

is required in order for covert rehearsal to succeed in maintaining a stimulus sequence in STM and to aid its transfer to long-term memory. For example, musical melodies are doubtless stored in an abstract (relational) symbol structure bearing no relation to the verbal system; but we do have a motor output system—humming or silent whistling—that allows us to regenerate a once-heard melody, and presumably thereby store it in memory better. But is there a similar covert rehearsal mechanism for odor sequences, for dance steps, or for sequences of tastes and touch patterns?

A second problem with the rehearsal buffer notion is that it assigns far too much weight to "dull, dead repetition" of the items themselves as the sine qua non for entry of items into more permanent memory. But subvocal naming of the items may only be the observable tip of an "iceberg" of deeper semantic processing. Specifically it now looks as if the transfer of information about items to LTM requires that those new items be related to familiar things in LTM and that interrelations among the new items be noticed and used for learning.

This viewpoint has been suggested by several considerations (see Craik, 1973; Jacoby, 1973; Mandler, 1967). First are increasing reports of situations in which verbal rehearsal of materials has virtually no influence on long-term retention of that material. Craik (1973) has distinguished Type I (superficial maintenance) rehearsal from Type II (semantic, organizational) rehearsal. Type I rehearsal causes verbal items to be maintained for a brief time in phonetic recirculation, but it results in very little storage of those items in LTM. Type II rehearsal involves more activation and consideration of the semantic meanings and associations of the items, trying to interrelate several of the to-be-learned items together in a meaningful sentence, image, or the like. An elementary experiment by Jacoby (1973) will illustrate the distinction. Adults listened to five unrelated words (well below their memory span) and either recalled them immediately or spent 15 sec rehearsing them subvocally before recalling them. Though recall was perfect in either case, items receiving the 15 sec of rehearsal certainly had a much higher frequency and duration of residence in the STM. After doing a number of such short lists, Jacoby surprised his subjects by asking them at the end of the session to recall any and all the items they could from the entire experiment. Significantly, recall of the multiply rehearsed items was no better than recall of the immediately recalled items—a clear failure of effect of repetition. On the other hand, we know that if, during that 15-sec interval, the subject had been asked to organize the five list words into a story or into an interactive scene in imagery, then later recall would have been greatly enhanced. A number of such demonstrations have shown us that people have a way of temporarily holding information in a phonetic form that produces almost no longer-term memory for the material. The distinguishing mark of experiments illustrating this point is that

the subject is not expecting the long-term retention test; he is, in fact, led to believe that once the material is "dumped from STM," then it can be forgotten. So it would appear that maintenance of items in STM is not a *sufficient* condition for the transfer of much information about them to LTS.

The converse side of the coin is that "residence time" of an item in STM is not so strong a predictor of long-term retention as is the type of processing the item receives during its residence in STM. Again, recent experiments by Craik (1973) and Jacoby (1973) prove this point as did earlier demonstrations by Mandler (1967), Tulving (1966), and Bower and Winzenz (1970). In Craik's experiment, the subject saw a word briefly after being set to judge it either in terms of its physical appearance (upper- or lower-case letters), its sound (rhymes with *hat*), its semantic category (is an "animal"), or its fittingness in a sentence frame ("blind as a _____"). Later free recall and recognition indicated much better memory for words that had been processed to the "deeper" semantic levels. An explanation for these results is currently being pursued. A likely explanation for the recognition memory results would begin by noting that recognition memory depends upon retrieval by the test item of information about the context of study of that item. We may suppose that the test word is better able to produce recall of its study question following the semantic-classification task simply because the word-to-category link has been primed by the study question and it preexists in memory, is strong, and is unique, whereas the word-to-type-font link (for the "shallow" processing) is novel, weak, and interfered with by many other items (for example, many items appear in upper case letters). Similar factors would imply better free recall for semantic than for case-queried items if it is supposed that semantic categories serve as the normal retrieval cues for free recall (and these *were* unique cues in Craik's experiment), whereas the type font or the sound of the item is not a self-generated retrieval cue. If this line of explanation were to be successfully extended, then a plausible interpretation of the "depth of processing" result is that it reflects the retrievability of the "question-to-item" event, which in turn depends upon the way the question biases the encoding of the item and the quality (uniqueness, prior associative connection) of the question-to-item association. These are currently active items on the research agenda.

C. Working Memory

To continue with the system architecture of Figure 2, we use the term working memory (or intermediate-term memory, ITM) to refer to the memory structures which maintain information about the local context, but information which is neither in the focus of active memory nor in the

distant edges of long-term memory. One of the primary functions of working memory is to build up and maintain an internal model of the immediate environment and what has been happening in our world over the past minute or two. We may think of the working memory as containing a description of the *setting,* framework, or context within which the more dynamic alternations of the world before us are taking place.

Several arguments indicate that we must have these temporary "world frames." First, visual perception is itself fleeting, skitterish, darting hither and yonder about a scene. But yet we do not every 200 msec recognize and construct the entire scene anew. Rather we integrate the information from those glances into one sustained image or model of the scene "out there" and our place in it. The model tells us the objects out there and their distribution in space with respect to ourselves.

This local model serves as a framework within which dynamic (small) changes are recorded—old objects are deleted, new ones are added, things change their relative positions, etc. But only the new, altered information must be focused on (in active STM) and entered as a new symbol structure or proposition into working memory. That is, new information "updates" rather than completely casting aside the current model.

These temporary models serve as a context for perception. In doing so, they constantly call upon information stored in long-term memory concerning the nature of objects and spatial relationships. For example, if I have just walked past a man who is climbing into his car, I am not startled to hear behind me the noise of his motor starting up a few seconds later. We "perceive" the rumble as a car motor rather than thunder or a building collapsing because our contextual model makes "motor noises" plausible. It is on the basis of long-term memory that I know that the lamp shade suspended apparently in midair across the room from me is in fact supported by that steel rod disappearing beneath it, and that a wire from the rod to the wall plug carries electric current to the bulb. So I am prepared to see certain things if I go over there and look under the lamp shade; also I expect certain changes in my "world model" if I pull out the electric plug or cut the wire. The point of such elementary examples is to say that perception is very much a matter of "predicting" the results of my interactions with my environment, and it is my current world model (based on partial sensory information) that enables me to bring the appropriate knowledge to bear on this issue.

Another phenomenon illustrating the need for an intermediate-term memory is in keeping track of what topics and referents have been recently mentioned during linguistic discourse. This is required in order to find the appropriate referents of anaphoric and pronominal forms. Also, it is used in finding the correct sense of ambiguous expressions. If I mention that

"John is a baker," later in the conversation I can refer to him as "the baker" or simply as "he,"; but to do so requires the listener (and me) to have a memory that keeps track of who is under consideration. Similarly, if the discourse is about a birthday party for little Billy, we must have that theme carried along in working memory to activate belief structures which provide proper interpretations and reasons for ambiguous sentences like "I'll buy him a top" or pronoun reference as in, "Don't give him one. He already has a top. If you do, he'll make you take it back." In this last example, the problem is to figure out that the *it* does not refer to the top that Billy already has but to the hypothetical "one" you are considering giving to him (see Charniak, 1972, for many examples). To summarize the point here, a working memory is necessary to keep track of foregrounded contexts in dialogues.

Another plausible use of a working memory is for it to hold proposed deviations from our actual world model that arise as we are thinking and trying to *imagine* the consequences of making certain alterations in our local environment. For instance, while looking at a configuration of chessmen, the master will try out (in imagination) a sequence of moves, entering the alterations in board positions of the relevant chessman in his temporary working model of the board, and then evaluating these projected imaginary states to decide upon a move. Similarly, an interior decorator rearranging furniture in a living room can use his working memory as a temporary imaginal model ("sketch pad") for trying out and evaluating various arrangements.

Finally, laboratory psychologists might identify the contents of the subject's working memory with the full set of propositions he knows about the experimental context in which he finds himself while engaged in, say, a nonsense-syllable learning experiment. At any moment we could ask him, "Where are you and what have you been doing?," and (with eyes closed) he could relate his complete spatial orientation to us, the room layout, where the materials-to-be-learned are displayed, what types they are, etc. He could also tell us what his instructions are, his strategies for dealing with the learning task, and how he evaluates the task as well as his performance. All of this would serve as an immediate mnemonic background or context within which his active short-term memory has just recorded episodic propositions like, "most recent letters in the memory-drum window were X–then–P–then–H–then–R." Through instructions, the subject knows that when we ask for recall, he is in fact to enter the retrieval cue, "most recent letters in the memory-drum window were ???." To use our earlier terminology, the working memory contains (among other things) the *plans* that the subject has initiated for dealing with the task we have set for him.

D. Long-Term Memory

1. General Considerations

The long-term memory (LTM) is assumed to be the "repository" of our more permanent knowledge and skills. It essentially includes all things that are in memory that are not currently being used. The information structures in LTM includes such general classes as the following:

1. Our spatial model of the world surrounding us—symbol structures corresponding to images of our house, city, country, and planet, and information about where significant objects are located in that cognitive map.
2. Our knowledge of physical laws, cosmology, of the properties of objects and things.
3. Our beliefs about people, about ourselves, about how to behave in various social situations. Our values and the social goals that we seek.
4. Our motor skills for driving, bicycling, shooting pool, etc. Our problem-solving skills for various domains. Our plans for how to achieve various things.
5. Our perceptual skills in understanding language or interpreting paintings or music.

Even at such a general level, the listing is hardly complete. Some of the information structures, when activated, produce a conscious experience of sensory imagery, some produce only a verbal flow (words of a memorized poem); others produce no introspective feedback whatsoever (for example, the "rules" we appear to follow in understanding speech).

At present there are just the beginning of a few models of LTM being proposed, specifically those by Anderson and Bower (1973), Kintsch (1974), and Norman and Rumelhart (1975). To these may be added the efforts of many computational linguists (e.g., Quillian, Schank, Winograd) who have worked on models for language understanding; such models require a highly organized memory, one that interacts continuously with a syntax parser in arriving at an internal description of an incoming sentence. The latter have been concerned with constructing memories suited for language parsing, but have not addressed themselves specifically to issues of learning, retrieval, forgetting, and such matters. All the systems are, in fact, oriented toward dealing with *sentence* inputs—with propositional "stimuli"—and they discuss other matters only by metaphorical extension of the way their systems deal with English sentences. All of these theories assume that LTM is a conceptual network, which serves as the "data base"

into which information (propositions) is inserted through the working memory and from which the central processor retrieves answers to questions. Conceptual networks were hit upon as a way to implement some general properties of human memory; therefore, to explicate such network theories, let us consider some of these properties.

2. Some General Properties of Human Memory

We list a few characteristics of the memory required and then describe a few functions that we would want a memory model to perform.

a. Associative retrieval. A first property is that information in memory should be retrieved in response to information that is associated with the input stimulus. To use a metaphor, the location of information in memory is specified in part by the contents of the item itself. Thus, a given symbol description (e.g., GLURK) will gain access to a spot in LTM at which would be collected whatever information (if any) that the person had associated to that stimulus pattern. The EPAM model of Feigenbaum (1963) is one realization of a content-addressable system. (In EPAM, the memory location corresponding to a stimulus item is arrived at by sorting the description of the stimulus through a series of feature tests.) Computer memories are not organized in this manner: rather data is placed seriatum in cells having locational addresses, but the address is unrelated to the datum stored there. An associative or "content-addressing" system can be approximated on current computers by "hash coding"; in such schemes the memory address containing information about an item is computed as some "hash function" of the item's contents.

The important issue for models of human memory concerns the units level at which the system is content addressable. Models like HAM assume direct access to sensory features and sensory relations; but given a perceptual description of a stimulus object, the system makes use of associative retrieval from there onwards. That is, a sensory feature pattern might "associatively retrieve" a word, and then a compound semantic idea (for example, horse meat, golf ball).

b. Concepts as nodes. The basic elements of memory are concepts (symbols) and relations between concepts (symbol structures). A concept may be a perceptual primitive (for example, vertical line, blue, horizontal movement), an actional primitive (for example, move hand, grasp, chew, suck, release), a primitive relation (for example, inside, beside, part of, has as parts, is a member of, above), or a higher-order concept built up by relations among these primitives. The concepts can stand for both generic terms (cups, pencils, governments) as well as individual constants (specific people and places). The "meaning" of a concept is given partly by

the configuration of its relationships to other concepts and partly by the referential conditions necessary for the proper use and application of the term. In our representation and in diagrams, concepts are represented by nodes (or cells in a computer memory) and relations between concepts as arcs, arrows, or labeled associations between the nodes.

c. Information as new conceptual configurations. In such a conceptual network, the learning of a new fact or new concept is solely a matter of recording its representation, which is in terms of a specified configuration of relations among already known concepts. Thus, for example, the concept of a helicopter would be introduced by listing its general type (airplane) and function, enumerating its distinguishing features (for example, propeller on top, capability for vertical takeoffs and landings), perhaps storing something equivalent to an "image" of the object. A new name may be associated to this concept, but the name is itself just another symbol structure (constructed from letters or phonemes). Most of our concepts do not have proper names but only definite descriptions (for example, "the person who stole my bicycle"), but they nonetheless would be represented in memory as a single symbol designating an entire symbol structure.

d. Concatenation principles. Our memories must have the capability to build up arbitrarily complex concepts from simpler concepts by use of simple rules of combination and concatenation. To do this, we must be able to say something further about any concept—that is, to attach another "relation-plus-concept" to any arbitrary concept. If I have the concept of the dog instance, Fido, I can predicate of him that he bit Harry, that I dislike him, that Jane owns him, and that Jane dislikes the fact that I dislike Fido. In these latter cases, notice that a *proposition* is itself being taken as an object (a "concept"). By representing propositions as nodes in memory, we are able to predicate about them as if they were simple concepts. From the viewpoint of the central processor, this means that a known proposition can be represented in working memory by a single token of its corresponding node in LTM. For example, we can deny a proposition by predicating "is false" of the proposition node.

e. Creating tokens from types. Any given concept, feature, or attribute will be used in many contexts, in encoding a great many events. For example, every time we store a proposition describing an event involving a new instance of a dog, we want to be able to create a new concept (the dog-instance) and say something about the event in which he occurs. We cannot predicate this particular event about the class of dogs; rather we must create in memory a *token* or *new instance* of the class of dogs, and attach the predication to this new instance. In the Anderson and Bower (1973) theory, this type–token link is simply one of set membership: Node

3714 (say) in memory is created to stand for a new member of the class of dogs, and we can then predicate many things about this dog instance (for example, its appearance and episodes in which it occurs). This capability for creating tokens of preexisting types is in fact the same capability noted earlier, of being able to create new concepts (for example, dog-3714) in terms of their relations to old concepts (for example, "dog–3714" is a dog).

f. Types of conceptualizations. The current models like HAM assume that the input to LTM is in a form equivalent to a proposition or a conceptualization. That is, LTM is capable of recording only inputs which have been "propositionalized" in an appropriate manner by the perceptual system. But this allows fantastically wide latitude. In the Anderson and Bower theory, a proposition is a subject–predicate construction put together according to certain rules. But such stipulations only invite one to supply the semantics of permissable combinations. Among the obvious types of semantic combinations which qualify as conceptualizations would be the following (suggested by Schank, 1973a):

> ACTORS perform ACTIONS
> ACTIONS have OBJECTS
> ACTIONS have INSTRUMENTS
> ACTIONS have RECIPIENTS
> ACTIONS may have DIRECTIONS
> OBJECTS can relate to other OBJECTS
>   these relations are: POSSESSION
>   LOCATION
>   CONTAINMENT
> OBJECTS can have ATTRIBUTES
> ACTIONS can have ATTRIBUTES
> ATTRIBUTES have VALUES
> CONCEPTUALIZATIONS can have TIMES
> CONCEPTUALIZATIONS can have LOCATIONS
> CONCEPTUALIZATIONS can CAUSE OBJECT'S ATTRIBUTES TO CHANGE VALUE
> CONCEPTUALIZATIONS can ENABLE other CONCEPTUALIZATIONS to occur
> CONCEPTUALIZATIONS can serve as REASONS for other CONCEPTUALIZATIONS [p. 5].

These presumably serve as some of the "syntax rules" for composing allowable conceptualizations. During conceptual development, particular concepts come to be classified as actors or objects, and particular actions come to be classified according to what conceptual cases they require (for example, "*physically give*" requires an actor, object, recipient, and direction of transfer of possession). Schank proposes that starting from a few primitive semantic notions and such rules as above, the infant acquires through perception a fragmentary conceptual knowledge of its environment onto which the structures of spoken language are then grafted. Then the two systems of language and world knowledge bootstrap each other up to the sophisticated cognitive machine that we are.
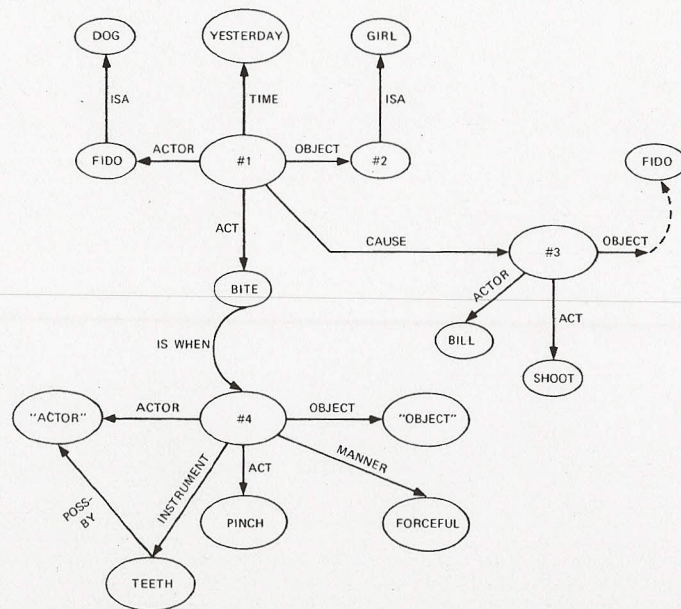
FIGURE 7   Diagram of the associative structure set up in memory to encode the conceptual content of "Fido bit a girl yesterday, so Bill shot him." The arrows are labeled associations (relations). The definition of *bite* is also shown. Definitions and knowledge about each concept is not shown to reduce clutter.

Putting aside the validity of such conjectures (or indeed, the more basic question of how to test them), it is fair to say that the memory network theorists seek a set of basic concepts and combination rules which will allow their system to represent practically any statement, event, or state of affairs. In particular, happenings or *events* in the world are to be represented in terms of a new configuration of relationships between tokens of preexisting concepts. To illustrate, look at Figure 7, which is a diagram representing the event "Fido bit a girl yesterday, so Bill shot him." Concepts are represented by circles, and relations by arrows. Node 1 represents the first event of Fido biting a girl-instance yesterday, and this event is said to cause the event in Node 3, that Bill shoots Fido. The small piece of network beneath *bite* illustrates the type of lexical decomposition into meaning components that linguists propose for most verbs. Thus, *bite* means for the actor to forcefully pinch the object using his teeth.

Such representations allow for retrieval of information by questions (probes) which stipulate part of the memory structure. Questions act like compound "retrieval cues," to use terms familiar in the memory field. If we ask, "Who did Fido bite?," the entry points (Fido, bite) retrieve Node 1 from memory; the interpreter then asks for "OBJECT-link of Node 1," which returns with Node 2; Node 2 is then found to be a girl instance—and

that determines the answer, "Fido bit a girl." If we ask "*Why* did Bill shoot Fido?," the system looks up the conceptualization of the query, getting a match to Node 3; it then backs up the causal link to Proposition 1, which permits it to construct the answer, "because Fido bit a girl yesterday."

g. Flexible retrieval. Another property we require of a humanoid memory is that a given memory be accessible by means of a variety of different cues. The supposed advantage of encoding events into *conceptual* representations rather than literal word strings or pictures is that the issue of flexibility of input or output by paraphrasing is handled in the input parser or the output generator (which are peripheral to the memory itself). Thus, I might have *seen* Fido bite the girl, and later learn verbally that because of that event Bill shot him; but my core memory representation is supposedly about the same as above. Similarly I would get the same memory structure if I had heard "Fido was shot by Bill because he found out that Fido had bitten a girl yesterday," or some other such paraphrase of the atomic propositions. Conceptual networks also have the potential flexibility for retrieval by paraphrases. Thus, I could ask, "Did a dog bite a girl?" requiring simple looking up of set membership for Fido; or I can ask, "Who forcefully pinched whom yesterday with his teeth?" and have hopes of perhaps getting a sensible matching answer at the level of concepts.

h. Ability to evaluate inputs. Another basic trait of human memory is that we constantly evaluate our perceptual or linguistic inputs to check for their consistency with what we know. We evaluate the truth of what is told to us; we answer "true" or "un-hm" if an input statement matches its conceptual counterpart in memory, and answer "false" or simply balk if we come across contradictory information ("Fido bit a boy yesterday"). Similarly, our memory provides an almost immediate decision that an input violates syntactic rules or semantic selectional restrictions ("Fido bit a ritual yesterday"). The exact means by which these contradictions and plausibility judgments are arrived at is a topic of continuing concern in research on semantic question answering.

i. Inference potential. It has long been recognized that many of the questions humans answer (read "much of the information retrieved from memory") is based on inferences from what we know rather than simply being a direct reading out of stored propositions. One way to view this is to think of the facts stored in our memory as like "axioms" of a logical system, so that questions are answered (if not directly matched) by applying a set of "subjective derivation rules" to our knowledge to come up with the answer. Thus, we derive that ostriches have lungs because they

live above ground and all such animals have lungs; or derive that Spinoza had an elbow because he was a man and all men have elbows. Such examples use the logical properties of relations like "set membership" and "has as parts"; that is, if x is a member of set y, and all y's have property p, then x has property p. A large number of relations (for example, "is located at") have some such logical properties.

A current focus of theoretical and experimental activity concerns how such inferences are made in semantic memory, particularly for set-membership statements ("a crocodile is a reptile") and for simple "has property" statements ("a crocodile has a stomach"). The competing theories must take account of the quantifier of the statement ("all, many, some, few, or no trees have leaves"), whether or not the subject and predicate are closely related in meaning, and, of course, whether the statement is true or false (for a good example, see Glass & Holyoak, in press; Glass, Holyoak, & O'Dell, 1974). Theories differ according to their assumptions regarding the organization of our conceptual knowledge and the search mechanism that operates over that network. Though the best model for the domain is in dispute, progress is clearly being made in research on the topic.

The discussion above refers to inferences drawn from long-term semantic memory. But another salient feature of humans is that they use a current event, episode, or statement to set off a chain of inferences about the context and situation described by the statement—inferences about the probable antecedents or reasons for the event, the probable consequences of it, and so on. We fit the event into our "model of the world," and are then able to expand it by inferences about various consequences. Consider just the unfortunate Fido incident above. We may infer any and all of the following:

1. Fido has teeth, probably sharp, and he has jaws.
2. His teeth made physical contact with some body part of a girl.
3. This probably caused the girl to feel pain and to cry.
4. Bill learned of this event.
5. It made Bill angry at Fido.
6. He decided to dispose of Fido.
7. He came in possession of a gun.
8. The gun was loaded.
9. Bill aimed it at Fido and pulled the trigger.
10. A high-velocity bullet passed through Fido's body.
11. This probably caused Fido to become dead.

This is but a brief listing of the myriad relations an average listener can "derive" or infer from the sentence said several pages back. The derivations are obviously of several different kinds: some come just from the meaning

of the action verb, some from the preconditions necessary for the given action to occur (for example, Fido and the girl had to be in the same location); other inferences come from our layman's postulates about emotional or causal consequences of actions and our reasons for actions. Abelson (1973) has written cogently on these issues, and Rieger (1974) has developed a simulation program that generates a large network of such inferences surrounding a given input sentence. Rieger argues that such inferences are frequently needed to follow the reasoning that weaves itself throughout even the most simple conversation. For example, if I know that Bill does not own a gun, then when I hear the original sentence the inferences stumble upon a "gap," which leads me to ask something like, "How did Bill get a gun?"

A point glossed over in the above illustration is that we have introduced "cause" as a relation between two propositions—between two events or states. Other propositional connectives would be *and, then, while, or,* etc. From such connectives one can represent stretches of coherent text much longer than single sentences. Thus, one can imagine giving a hierarchical description to stories. Speaking abstractly, a given story might consist of, say, Episode 1, then Episode 2, which cause Episode 3. But Episode 1 in turn consists of Event 1 together with Event 2, both causing Event 3, which completes Episode 1, and so initiates Episode 2 of the story, etc. Analysis of the conceptual structure of texts and stories, and their encoding into memory, is just beginning (see Crothers, 1972; Frederiksen, 1972; Meyer, 1974; Rumelhart, 1973), but this is clearly going to be a very important area for future research.

A current topic of interest to linguists is the issue of how to classify and decompose verbs into a small set of primary meaning components. Schank (1972, 1973b) and Norman and Rumelhart (1975) have devoted major efforts to this topic. The motivating philosophy is much like that behind early chemists' search for basic atoms from which complex molecules would be composed. One implication, for example, is that if different meaning components become available at different stages during cognitive development, then in learning his language a child should progress through an orderly set of "stages" at which he understands progressively more complex verbs (for example, *have,* then *give,* then *sell*) while being confused among verbs that use semantic distinctions he has not yet acquired.

3. Learning in Conceptual Networks

The treatment of cognitive learning in these network models is perhaps obvious given the foregoing background. When a new fact (proposition, statement, scene) is entered into working memory, we suppose that it corresponds to a structured conceptual tree diagram such as shown in Figure

7. What is "in working memory" at this time are the tokens of the LTM concepts and the temporary connections (labeled links, associations) that are to be established in order to remember the fact. Learning consists in recording these linkages into long-term memory.

An initial question is, "When will any learning be initiated?" This would seem to depend on the "informativeness" (or unpredictability) of the stimulus sequence or event under consideration. We assumed earlier that each incoming proposition is first checked for whether all or part of it is already known (is stored in LTM). If it is a familiar fact, then little is stored about it—perhaps only the further fact that it was mentioned in the present context. Since working memory is constantly using part of the current world model and LTM to predict forth-coming events (for example, if lightning, expect thunder), the foregoing assumptions indicate that the occurrence of predicted events is easily perceived but unlikely to activate any efforts toward new learning of coincidental events.

If an incomng statement finds a partial match between its information and an existing tree structure in LTM, then that matching structure is used as the node onto which the difference information ("new predications") are attached. If I already have stored that Fido bit a girl and I am later told that the girl Fido bit was little, rich, and blonde, I do not rerecord all the nodes and links encoding the original fact. I simply add to that structure (specifically, to Node 2 in Figure 7) the new attributes asserted about the girl instance. The important assumption here, brought out particularly in the learning routines of the Anderson–Bower theory, is that, whenever possible, incoming information is recoded and recorded by re-using and modifying parts of already known information structures. This is economical in eliminating redundancies of a certain sort. It also illustrates why highly associated material is so easily learned (for example, compare the learning of "the doctor cured the patient" versus "the doctor slapped the hostage")—because most or all of it is already known. The hypothesis also accommodates our tendency to *assimilate* new information in terms of familiar conventions and stereotypes (see Bartlett, 1932). We may store an event as "Event 7 has familiar Description 172 except for Differences 1, 3, 6." But then we forget the differences (they were not successfully recorded in LTM) and so reconstruct our memory of the event or story in familiar, conventional form.

Let us consider finally the case where the incoming statement expresses totally new propositions—they are sensible but not already known in LTM. How then does learning proceed? Presumably by rehearsal of the network of relations among the concepts and by thinking about and perhaps "imaging" the referents of the concepts and their stipulated interrelations. As a *quantitative* description, it may be assumed that in each small unit of time as the person is concentrating on the proposition, there is a fixed

probability that any given link in the input tree will be recorded into long-term memory. Given a fixed study time per sentence (typically, 3–10 sec in most learning experiments), the the number of relational links that are durably formed in a tree is a probabilistic matter. That is, the full statement is not learned all or none as a unit; rather some subparts may be learned before others (see Anderson & Bower, 1973, Chapter 10).

The result of this stochastic learning assumption is that there will be many "fragmentary trees" in memory that result in partial, vague memories. And such fragmentary memories do indeed exist, either when we introspect or examine experimental recall results. Thus, you might remember that "Fido bit someone, for which he was shot," but forget who was bitten, when it occurred, and who shot him. Such fragments, plus inferential elaborations done at the time of learning or recall, provide the material for plausible, imaginative reconstructions done according to normative knowledge. For example, I might elaborate my fragments to "Fido probably had rabies, was running wild in the park where he bit an old man, and so was shot by a passing policeman who witnessed the attack." Our conventional knowledge and stereotypes have elbow room to intrude in just such cases.

The propositional learning model given above is hardly profound, but there is empirical evidence in its favor (see Chapter 10 and 11 of Anderson & Bower, 1973) and no intuitively compelling alternative has been proposed. A problem with it is that it seems testable only given the naive assumption that a word given alone as a retrieval cue (for example, *bite*) is interpreted in the same semantic sense as it was when it occurred in the context of the study sentence (as in "dog bite"). In fact, however, subjects may occasionally think of a different sense of *bite* than that encoded during input (for example, *bite* can also mean to beg from someone, or to sting, or to clip one's speech). But these are second-order deviations from the basic model of probabilistic link formation of learning (see R. Anderson & Ortony, in press, for counterarguments).

E. Decision Making

An organism is faced with a "decision" when he has two or more alternative courses of action for each of which he must evaluate and compare the expected consequences. In a simple experimental situation involving two stimulus–action–outcome alternatives, the person obviously tries to accumulate knowledge allowing him to predict the payoffs or penalties following each stimulus—action combination. He essentially learns such correlations by accumulating rehearsal on propositions such as "pressing the left key to the cue MIB results in five points." For a rational decision, this bit of knowledge would be needed along with learning of a complementary proposition such as "pressing the right key to the cue MIB results

in three points." Given these two bits of knowledge and a higher valuation on five than on three points, then presentation of MIB will retrieve the two propositions and the rule of "choose the more preferred outcome" will cause the central processor to issue the command "execute a left-key press." Such is a simple matter.

More complex decisions involve either choices between two large commodity bundles (for example, a vacation trip to Yucatan versus Costa Meralda) or several outcome bundles having a probability distribution over the several outcomes (that is, bets). The approach of cognitive psychologists to the study of such decision making is basically descriptive rather than normative. For example, in studying how people evaluate bets, the psychologist will examine how decisions are affected by the importance different individuals assign to amounts of gains or losses, and to the probabilities of wins and losses (see Slovic & Lichtenstein, 1968). Whereas the evaluation of a single "outcome-by-probability" package seems predictable by multiplying the subjective scale values of the outcome and the probability, the total value of two such outcome-by-probability packages is consistently less than the sum of the values of the two packages taken separately (Shanteau, 1974). Of great interest to studies of decision making are recent papers by Kahneman and Tversky (1972, 1973) and Tversky and Kahneman (1971, 1973). They find that when people estimate probabilities of events, they use intuitive heuristics which produce systematic biases in their estimates. One heuristic is "representativeness," wherein people assign higher probability to any outcome that conforms more closely to their stereotype of what a "representative example" ought to look like. Thus, the sequence HTHT will be judged as a more probable outcome of four tosses of a fair coin than will the sequence HHHH, although objectively the two sequences are equally probable. A second heuristic biasing relative frequency judgments is relative availability in memory of the events being judged. Thus, for example, a person will readily generate more words beginning with the letter k than he will words having k in their third position (a fact about memory organization); correspondingly, he will also estimate a higher frequency for the former class of words—which, in fact, is false. The literature cited may be read to see some of the specific biases these heuristics cause when they are applied to various real-world instances of intuitive explanations and predictions.

F. Response Generation

Once an action has been decided on (such as to open the flue of the fireplace), a central plan is drawn from long-term memory for that action and is initialized (or "parameterized") for the current situation in which

the response is to be executed. It appears that most skilled actions may again be analyzed in terms of a multilevel hierarchy with various tests in the plan which compare accomplished with anticipated results. Interestingly, just as we may think of perception as the fitting of a concept onto the current sensory field, so may we think of action in terms of the subject trying to realize a certain concept in his own behavior. Subjects can obviously categorize others' behaviors according to concepts (for example, "he's striking a match," "he's writing his name"). It seems reasonable to suppose that, given an intention, the person then makes those movements necessary under the circumstances to result in his own behavior instancing that concept. Thus, we might think of an action family as those sequences of movements which pass the criteria for being classified under the circumstances as exemplars of a given action concept.

Most contemporary theories of motor organization rely heavily upon concepts from servocontrol theory (see Bernstein, 1967; Pew, 1974). Such theories assume the existence of components necessary for servoregulation of movements: *effectors* that move limbs, *sensors* that provide feedback, a *control element* that contains the current goal stipulated in terms of the desired setting of an intensity or location parameter, a *comparator* that determines the size and sign of the discrepancy between the target value and the achieved value, and a *translator* that converts a given discrepancy into appropriate correctional movements designed to reduce the discrepancy. For instance, in tracking an erratically oscillating spot of light (say, on a TV screen) with one's fingertip, the goal is to maintain a null discrepancy between the momentary location of the spot and the fingertip. To the extent that there are large delays between sensing a discrepancy and firing off a corrective movement the response of the system will lag (be in error) because the dot has moved elsewhere during the delay. The error can be reduced considerably if the movement of the dot is regular and periodic (for example, traces out a sine wave as it sweeps across the TV screen). A higher-level "anticipator" that has learned the periodicity can provide anticipatory commands to the control element moment by moment, so that the fingertip will be moved to where the dot is expected to be in the next instant. It is such predictive knowledge that allows the motor system to move beyond the low-level actions of merely corrective reactions. It is generally supposed that at the higher levels, the solution to a motor problem ("open the flue") is represented in advance in the brain as an abstract motor image of the person's surroundings, the actions to be carried out, the instruments and support relations to be used, and the expected effects of these actions. Such motor programs seem to be using again the basic ideas of hierarchical control: set a goal, select a plan or method, operate for awhile, check your results to see how you are progressing, and so on.
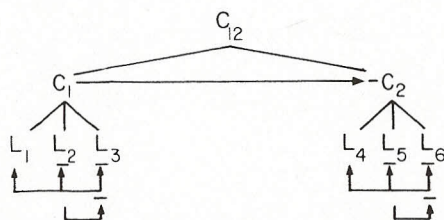
FIGURE 8 Diagram of the associative structure representing a two-chunk series, each composed of three letters. Lines indicate excitatory connections; arrows ending in a minus sign denote inhibitory connections. (Reproduced by permission from Estes, 1972.)

Along somewhat different lines, Johnson (1970) and Estes (1972) have been developing hierarchical models for the organization of serial responses. The typical domain of application is to a human subject learning to smoothly speak a chunked letter series such as HPJ–MXK (see also Lesgold & Bower, 1970). Estes provides an associative theory wherein higher-order "control elements" are introduced to represent chunks, with associations from this element to the items within the chunk. The order of elements is encoded in terms of *inhibitory* links between elements within a chunk. Figure 8 from Estes (1972) shows the internal representation of the series. The C's stand for control elements, the L's stand for the motor programs for uttering the particular letters, lines denote excitatory connections, and arrows ending in "minus signs" denote inhibitory connections. The way the structure fires off is as follows: the decision to respond causes activation of the top control element (C_{12}), so excitation then flows to its next-level elements (C_1 and C_2). But because C_1 inhibits C_2, C_1 will have a greater activation. It passes excitation to its letter tokens L_1, L_2, L_3, of which L_1 occurs first because it is least inhibited. After L_1 occurs, it releases its inhibition of L_2 and L_3. Then L_2 occurs, since it is more activated because it is still inhibiting L_3. After L_2, then L_3 occurs. This completes the chunk C_1, releasing inhibition on chunk C_2, so that it begins popping out L_4, L_5, then L_6. Estes' model for output is much the same as Johnson's, except for the conjecture regarding *inhibitory* connections (and their release) as encoding serial order.

One advantage of such models, however formulated, is that they can permit very rapid sequential responses which are organized centrally and are not strictly dependent upon moment-by-moment feedback stimulation providing the stimulus for the next response in the chain. In a critique of the early stimulus–response theory of chaining (that is, where feedback from response n was assumed to be the stimulus eliciting response $n + 1$ in the series), Lashley (1951) had objected that this circular mechanism could not operate fast enough to explain very fast sequential responses (for example, a skilled pianist racing through an arpeggio). Chunking models like Estes' solve the problem since all responses within a chunk are activated at once, and each can occur very rapidly as soon as the inhibitory input preventing its occurrence has been released by the evocation of its predecessor.

Although there is very much more that could be said about motoric processes—their relation to early knowledge representations (Bruner), to early imagery (Piaget), to thinking (Bartlett), and how speech is produced—space limitations prevent our going further into the matter here.

IV. COMMENTS ON LEARNING PARADIGMS

A. Induction

It is fitting in a volume devoted to cognition and learning to end with some brief comments on how the cognitive system sketched in the foregoing applies to the various learning paradigms that have been investigated so intensively by psychologists. Roughly speaking, learning experiments may be viewed as settings which expose the subject to controlled opportunities for acquiring information. This information typically refers to the correlations (arranged by the experimenter) between environmental events or between the subject's actions and environmental outcomes.

The old associationists saw that the human brain was a powerful induction device—a correlation detector—and they explained that ability in terms of its more primitive ability to associate ideas (corresponding to sensations or actions) according to their contiguity in time or space. Thus, I associate my easy chair with my fireplace because they occur together in space; I associate striking a match with its lighting because they occur contiguously in time.

Within certain boundary conditions and with emendations and exceptions to be noted, such principles seems to explain the learning that occurs in standard situations. A first simple emendation concerns the historical empiricists' theory of "sensation." They subscribed to the view that the sensory field projected an unstructured mass of sensory elements that caused various features and properties to "light up" in the brain, and hence to become interassociated by virtue of their being experienced contiguously in time and space. That view appears too simplistic, however: rather, the sensory field seems to become organized almost immediately in terms of objects and their interrelations. Gestalt principles of perceptual organization are brought to bear in segregating the field into objects, describing properties as unified with their objects, and identifying objects mainly in relationship to one another. The input to memory can not be considered to be a disorderly mosaic of punctate sensory elements; rather, the perceptual system provides working memory with highly structured contents for storing in LTM.

B. Conditions for Learning

A second emendation is that neither temporal contiguity nor spatial contiguity of *objective* events (or properties) are necessary or sufficient conditions for establishing permanent associations. What is critical is that the two mental representations enter into a single propositional structure in STM; this is the common result of those circumstances which promote "belongingness" of the two events (see Thorndike, 1931). Contiguity without belonging creates only temporary connections which rapidly fade from STM.

Besides the facts about the role of belongingness in learning, other facts suggest that contiguity is not sufficient for learning. One is our failure to learn the millions of series and patterns of sensory events that we experience every day. Incidental learning is notoriously poor (for example, what letters go with what spaces in your telephone dial?). Thousands of bits of structured information enter STM every few hours; most is never used, some is used only temporarily to update a rapidly changing world model (for example, "What down is it in the football game?"), and only a small fraction of the information is ever consolidated into LTM, and that is often of an abstract sort (for example, "our football team won"). The critical circumstances promoting LTM consolidation of contiguous events are not fully known but the question has been intensively investigated. With a human subject, any signal or instruction that informs him that the foregoing information is important and valuable will induce him to rehearse and enrich in some way the symbol structure describing the events and their relations. Instructions in human learning experiments typically serve this role of identifying for the subject exactly what is to be noticed and rehearsed.

With lower animals (and sometimes people) we often must use biologically important stimuli (rewards or punishers) to signal the same message, viz., "this information is valuable, so learn it." Food to a hungry rat is a significant event that introduces a pause or "articulation joint" in an otherwise homogeneous flow of behavioral events. The reward stops the flow momentarily; it essentially signals the central processor to "rehearse the preceding contents of short-term memory." N. E. Miller (1963) called this consolidating rehearsal the "go mechanism," and showed how it could explain a number of findings regarding reward and learning.

The learning mechanism seems to become "switched on" mainly when environmental events do not confirm expectations—when they are surprising or informative. This seems an appropriate rule to follow if a system were to be designed to learn new information but to not waste space storing redundant, predictable information. This assumption regarding informativeness seems supported by the research on the "blocking" of learning

(Egger & Miller, 1962; Rescorla, 1972): in such experiments, if Cue A already predicts some important event (say, a painful shock), then trials pairing Cues A and B before shock will result in no conditioning of Cue B. It is as though since A predicts shock and that correlation continues during the $A + B$ trials, there is no new surprising information to be learned, so the B-to-shock association is not rehearsed.

The work of Garcia and Koelling (1966) (see Bolles, Chapter 7 of this volume) indicates that the stimuli that are selected out of the record of the past to be associated with an unconditioned reaction depends on what that reaction is. These investigators showed that when a rat becomes sick to his stomach, he associates that selectively with his memory of the most recent novel taste he has experienced, disregarding the myriad auditory and visual stimuli that have passed through STM during the interval between the novel taste and his becoming sick. The result illustrates a selective "mental contiguity" of obvious biological utility (for example, it is most likely that you get stomach sickness from some unfamiliar thing you ate).

Such results suggest that contiguity of events in objective time is not a necessary condition for associating them. Similar selective associations of events across a time span of irrelevant events are easily demonstrated with people. For example, if someone sets my house on fire, I will selectively remember any recent episode in which some enemy threatened to burn my property, and associate (inferentially) the arson with that enemy. As an experimental example, Jacoby (1974) asked some of his subjects to judge whether each successively presented word was in the same semantic category as the just preceding item of the series; other subjects judged whether each item was in the same category as *any* preceding item of the series. In cases where two items of a category (dog–giraffe in the "animals" category) were presented many items apart, the latter subjects showed very much higher recall of one item given the other item as a cue ("dog" given "giraffe"). What was clearly important during the learning trial was the revival of the earlier instance and joint rehearsal of the two items at the time the second item occurred. Such "mental contiguity" was far more important for association formation that was "objective contiguity" of the items. In the Jacoby experiment, instructions designated the terms to be put into mental contiguity; in the Garcia and Koelling experiment, genetic programming dictated the selectivity.

C. Cognitions or Habits?

A perhaps distinguishing feature of the cognitive approach is that it assumes that "what is learned" are conceptual associations encoding knowledge about the relationships among events in the learning situation; this

may be contrasted with the traditional doctrine that a stimulus–response habit is what is learned. Thus, a human adult being trained on an elementary "operant discrimination" problem could detect the relevant correlations and frame for himself "if–then" propositions expressing such contingencies as, "when the high-pitched tone sounds, if I press this lever, I get a dime about every tenth press; if the high-pitched tone isn't on, then I get nothing for pressing." The point of calling that a "conceptual structure" is simply to recognize the versatility of ways by which we may instill, activate, or modify that simple knowledge structure, and the versatility of effective actions the structure will mediate in altered circumstances. For one thing, a human adult could learn by simply being told the contingencies, or by observing another person being exposed to the contingencies. The experimenter may alter details of the exact response (for example, "use your left hand rather than your right"), or payoff in convertible tokens rather than dimes, or convey complete information about altered contingencies (for example, "there'll be no more dimes"), and the subject will quickly alter his actions appropriately. The subject not only presses the lever but can describe the contingencies to himself or to a friend in need of dimes. He becomes "aware of" the contingencies, and only then does he begin to behave appropriately (provided he values dimes and the experimenter's good will). Recent reviews of the critical evidence regarding the role of awareness in human classical and instrumental conditioning (Brewer, 1974; Dulany, 1968) indicate that people put through such procedures figure out the contingencies (either exactly or with a closely correlated hypothesis), and then respond according to their expectations and the outcomes they want to get. "Figuring out the contingencies" simply means becoming aware of or able to label the stimulus–action–outcome correlations programmed into the situation by the experimenter.

Although our examples above use the verbally competent adult, we may suppose that a similarly abstract conceptual proposition encoding the temporal sequence of significant events is learned by nonverbal humans and animals. To quote from Anderson and Bower (1973):

> What complex "abstract structure" is, say, a dog learning when a tone is consistently followed by a painful, noxious stimulus? There is no reason to deny a priori that dogs and other nonverbal organisms can form the equivalents of elementary propositions encoding the temporal sequence of significant events. To be absurdly concrete, suppose the dog acquires an associative structure expressible as the proposition: "In experimental situation S, a high-pitched tone is followed within a few seconds by a painful stimulus to my left front paw." If this proposition is combined with other general propositions, such as "If a limb is about to be injured, move it away to avoid injury," the dog could "derive" the command to flex its left paw when the tone sounds. The claim is that such conceptual propositions capture the flavor of the dog's behavior, especially in new situations. (This was Edward Tolman's main position.) Thus, it is known that the overt paw-flexion response need not occur during learning. Blockage of the neuromuscular junctures with a curare-like drug prevents all overt responses, yet the tone-shock association may be established under curare, and then later tested with positive results (i.e., conditioned paw withdrawals to the tone) after the drug has worn off. Also the tone–pain association institutes in the dog a totally different "repertoire of responses" to the tone in other situations. Thus, in a later appetitive instrumental conditioning situation, the tone will serve as a "conditioned emotional stimulus" suppressing appetitive behavior. The tone can also be used as a "punishing" event, supporting passive avoidance learning; also, its onset can initiate various kinds of previously learned escape behavior, and its termination can "reinforce" escape behavior. There is, thus, a diverse class of different behaviors affected by the tone-shock experience. Beyond this, there are probably other means to induce a similar memory structure in the dog. For example, the tone could be paired with other painful stimuli like burns, loud noises, nausea, or pinpricks delivered to the left paw.
>
> The tenor of these comments is that for even a nonverbal organism and for a singularly simple connection such as "tone–shock," we still require for its representation an intervening propositional structure which can account for multiple determinants, multiple contexts of retrieval of that information, and multiple varied "behavioral indices" related to that learning [pp. 36–37].

Although such considerations argue that the effects of a conditioning history may be represented in terms of a propositional structure, the information-processing approach itself does not claim to "solve" the many scientific puzzles that surround investigations of conditioning and learning. The theoretical language we use does not automatically tell us when conditioning will occur and when it will fail, how reinforcing events are to be characterized, how parameters of the learning situation exert their influence on the process, or how species-specific action patterns may intrude to override and interact with processes of associative learning. These are genuine puzzles to be investigated regardless of whether one adopts as his theoretical language stimulus–response concepts or information-processing concepts. With that general disclaimer, let us consider briefly the two issues of conditioning rate and forgetting.

D. Contextual Variability and Learning Rate

In analyzing conditioning, we have assumed that the unconditioned stimulus enters STM and also activates rehearsal of the contents of STM. Since these contents become associated to the unconditioned stimulus (US), why is conditioning not always complete in one trial? The answer is that the speed of learning will depend on the variability over trials in the contents of STM and working memory at the time the UCS arrives. Although standard experimental arrangements almost insure that the critical to-be-conditioned element (CS) is in the STM at the time of the US, one cannot

control and keep constant the sensory or ideational contents that are coincidentally filling STM at the start of the trial or during the interval between the CS and the US. This variability in contextual elements in STM acts very much like the trial-by-trial variability of stimulus samples as envisioned in the stimulus sampling theory of Estes (see Estes & Suppes, 1974). The contents of STM on a given trial early in learning will serve as a set of independent retrieval cues; since only a few of these cues will retrieve the prior association to the US, the US will be anticipated with only a low probability. Over trials, the critical elements in the CS–US correlation will be isolated due to variation of irrelevant details.

To illustrate, suppose that on Trial 1 the contents of STM are elements a, b, c, d when the UCS occurs; on Trial 2 they are a, b, x, y; on Trial 3 they are z, b, x, k; and so on. We would like to specify a learning algorithm that will eventually assign a high weight to element b in calculating an expectation of the US. A simple algorithm which does so is one that increases the weight (attentional strength or perceptual importance) of an element whenever it occurs with the event to be predicted (US), and which decreases the weight of the element either if (1) the stimulus element occurs without the US, or (2) the element occurs in the presence of a more valid predictor of the US, or (3) the US occurs in the absence of the stimulus element (see Rescorla, 1972). Such an algorithm essentially states a mechanical procedure for the brain to compute trial by trial the current correlation between presence or absence of a cue and presence or absence of the UCS. After learning, then, the most valid cue (the CS) tends to receive high attentional weighting and thus the subject's expectation will be determined by the association retrieved by this most valid cue in the STM during a trial.

We thus view Pavlovian conditioning essentially as "discrimination learning," as the subject learning to isolate the most valid predictive cue (of the US) in a variable environment. Operant conditioning would be viewed similarly except the important element entering STM would be feedback from (or the efferent command issued for) the response upon which reward is contingent. That is, the subject would learn a correlation of the form, "in situation S response R is followed by outcome X." Such knowledge modules are then retrieved and activated whenever the subject is in a situation resembling S and has a present demand or desire for outcome X.

E. Retention of Learned Associations

The primary factors contributing to retention losses of well established associations are contextual changes in the prevailing stimuli and possibly the learning of competing events or responses to the stimulus during the retention interval. The chances of these negative factors occurring increases over

the usual retention interval. Contextual changes refer to alterations over time in the setting and thoughts that fill STM or the direction of attention (to different sensory dimensions), so that an alered pattern of stimulation serves as the retrieval compound. The compound is less likely to retrieve the designated association the greater the contextual change (see Bower, 1972, or Estes, 1955, regarding the "fluctuation" model of contextual alterations and their effect on retention).

The second factor promoting forgetting, namely interfering learning, has been much studied, and the basic principles seem to apply quite generally to various paradigms, responses, and species (see Postman, 1971). If the subject learns an association A–B, then to the cue of A he is likely to forget it or to become confused in remembering if he has also learned the competing association A–C either *before* or *after* the A–B learning. Cognitive theory has no distinctive hypothesis to offer regarding the underlying causes of proactive or retroactive interference; they are simply accepted as phenomena in need of explanation. Anderson and Bower (1973, Chapter 15) offered the interpretation that the events A–B were acquired by elaboration of a proposition (say A as subject, B in the predicate), that A–C was acquired similarly, and that this results in a temporally ordered stack of associations out of A. An inherent restriction on retrieval times would then produce the observable phenomena of retroactive interference, whereas restricted retrieval plus a random reordering of the association stack out of A would produce the observable phenomena of proactive interference. The references cited should be checked for arguments for and against this hypothesis.

F. Cognitive Psychology and Social Learning Theory

While commenting briefly on the way cognitive psychology applies to some laboratory learning task, it is appropriate to note that cognitive psychology helps us also understand the learning and utilization of so-called "personality" and social behaviors. Its applications to socialization and personality development are at least as cogent, if not more so, than the account provided by S–R reinforcement theory. To cite but one example, there has been a continuing complaint that the so-called "behavior therapies" or "behavior modification procedures" use techniques of diagnosis, persuasion, motivation, imagery, observational modeling, thought control and image control, etc. in their therapy, despite the fact that none of these techniques are at all explicable in terms of the basic framework of behaviorism and conditioning (see Breger & McGaugh, 1965; Locke, 1971). The ascendant if not prevailing viewpoint today in accounts of socialization, self-control, and personality development is "social learning theory" (e.g., Bandura 1969, 1971; Mischel 1969, 1973). It differs from S–R reinforcement theory in its emphasis on mechanisms of observational learning,

self-control, and "cognitive factors" in the acquisition and utilization of social behaviors. Social learning theorists refer repeatedly to "internal representations" of observed situation–action–outcome sequences (for example, in language codes, in images, or in conceptual propositions), and how these are remembered, utilized, and guided by anticipation of the desired outcome. It seems that cognitive psychologists are studying those intellectual skills that social learning theorists feel compelled to refer to repeatedly in explaining how a person acquires knowledge of his social environment and acquires skill in negotiating his way through it. Thus, amalgamation of cognitive psychology with social learning theory should provide a scientific but broadly relevant synthesis of a sort that psychology has so long been searching for.

ACKNOWLEDGMENTS

REFERENCES

Abelson, R. P. The structure of belief systems. In R. C. Schank & K. M. Colby (Eds.), *Computer models of thought and language*. San Francisco: Freeman, 1973.

Anderson, J. R., & Bower, G. H. *Human associative memory*. Washington, D.C.; Winston, 1973.

Anderson, R., & Ortony, A. On putting apples into bottles: A problem of polysemy. *Cognitive Psychology,* in press.

Appel, L. F., Cooper, R. G., McCarrell, N., Sims-Knight, J., Yussen, S. R., & Flavell, J. H. The development of the distinction between perceiving and memorizing. *Child Development,* 1972, **43,** 1365–1381.

Atkinson, R. C., & Shiffrin, R. M. Human memory: a proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation.* Vol. 2. New York: Academic Press, 1968.

Atkinson, R. C. & Shiffrin, R. M. The control of short-term memory. *Scientific American,* August, 1971, 82–90.

Bandura, A. *Principles of behavior modification.* New York: Holt, Rinehart & Winston, 1969.

Bandura, A. *Social learning theory.* New York: General Learning Press, 1971.

Bartlett, F. C. *Remembering.* Cambridge, England: Cambridge University Press, 1932.

Belmont, J. M., & Butterfield, E. C. Learning strategies as determinants of memory deficiences. *Cognitive Psychology,* 1971, **2,** 411–420.

Bernstein, N. *The Co-ordination and regulation of movements.* Oxford: Pergamon Press, 1967.

Bjork, R. A. Theoretical implications of directed forgetting. In A. W. Melton & E. Martin (Eds.) *Coding processes in human memory.* Washington, D.C.: Winston, 1972.

Blakemore, C. Development of the mammalian visual system. *British Medical Bulletin,* 1974, **30**(No. 2), 152–157.

Bower, G. H. A descriptive theory of human memory. Paper given at the Second Annual Princeton Conference on Learning and Remembering, Princeton, New Jersey, 1964. Proceedings published in D. Kimble, (Ed.), *Learning, memory and forgetting.* Vol. 2, New York: New York Academy of Science, 1967.

Bower, G. H. Stimulus-sampling theory of encoding variability. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: Winston, 1972.

Bower, G. H., & Winzenz, D. Comparison of associative learning strategies. *Psychonomic Science,* 1970, **20,** 119–120.

Breger, L., & McGaugh, J. L. Critique and reformulation of "learning theory" approaches to psychotherapy and neurosis. *Psychological Bulletin,* 1965, **63,** 338–358.

Brewer, W. F. There is no convincing evidence for operant and classical conditioning in humans. In W. B. Weimer & D. S. Palermo (Eds.), *Cognition and Symbolic Processes.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1974.

Broadbest, D. E. *Perception and communication.* New York: Pergamon, 1958.

Charniak, E. Toward a model of children's story comprehension. Technical Report AI TR-266. Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1972.

Conrad, R. Acoustic confusions in immediate memory. *British Journal of Psychology,* 1964, **55,** 75–84.

Craik, F. I. M. A "levels of analysis" view of memory. In P. Pliner, L. Krames, & T. Alloway (Eds.), *Communication and affect: Language and thought.* New York: Academic Press, 1973.

Crothers, E. J. Memory structure and the recall of discourse. In J. B. Carroll and R. O. Freedle, (Eds.) *Language comprehension and the acquisition of knowledge.* Washington, D.C.: Winston, 1972.

Duda, R. O., & Hart, P. E. *Pattern classification and scene analysis.* New York: Wiley–Interscience, 1973.

Dulany, D. E., Jr. Awareness, rules and propositional control: a confrontation with S–R behavior theory. In D. Horton & T. Dixon (Eds.), *Verbal behavior and general behavior theory.* Englewood Cliffs, New Jersey: Prentice-Hall, 1968.

Egger, M. D., & Miller, N. E. Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology,* 1962, **64,** 97–104.

Ellis, N. R., Detterman, D. K., Runcie, D., McCarver, R. B., & Craig, E. N. Amnesic effects in short term memory. *Journal of Experimental Psychology* 1971, **89,** 357–361.

Epstein, W. Mechanisms of directed forgetting. In G. H. Bower (Ed.), *The psychology of learning and motivation.* Vol. 6. New York: Academic Press, 1972. Pp. 147–192.

Estes, W. K. Statistical theory of spontaneous recovery and regression. *Psychological Review,* 1955, **62,** 145–154.

Estes, W. K. Reinforcement in human learning. In J. Tapp (Ed.), *Reinforcement and behavior.* New York: Academic Press, 1969.

Estes, W. K. An associative basis for coding and organization. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: Winston, 1972. Pp. 161–190.

Estes, W. K. Phonetic coding and rehearsal in short-term memory for letter strings. *Journal of Verbal Learning and Verbal Behavior,* 1973, **12,** 360–372.

Estes, W. K., & Suppes, P. Foundations of stimulus sampling theory. In D. H. Krantz, R. C. Atkinson, R. D. Luce, & (Eds.), *Contemporary developments in mathemati-*

cal psychology. Vol. I: *Learning, memory, and thinking.* San Francisco, California: Freeman, 1974.

Feigenbaum, E. A. Stimulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman, (Eds.), *Computers and thought.* New York: McGraw-Hill, 1963.

Flavell, J. H., Friedrichs, A. G., & Hoyt, J. D. Developmental changes in memorization processes. *Cognitive Psychology,* 1970, **1,** 324–340.

Frederiksen, C. H. Effects of task-induced cognitive operations on comprehension and memory processes. In J. B. Carroll & R. O. Freedle (Eds.), *Language comprehension and acquisition.* Washington, D.C.: Winston, 1972.

Garcia, J., & Koelling, R. A. Relation of cue to consequence in avoidance learning. *Psychonomic Science,* 1966, **4,** 123–124.

Garner, W. R. *Uncertainty and structure as psychological concepts.* New York: Wiley, 1962.

Gibson, E. J. *Principles of perceptual learning and development.* New York: Appleton-Century-Crofts, 1969.

Glanzer, N., & Clark, W. H. Accuracy of perceptual recall: An analysis of organization. *Journal of Verbal Learning and Verbal Behavior,* 1963, **1,** 289–299.

Glass, A. L., & Holyoak, K. J. Alternative conceptions of semantic memory. *Cognition,* in press.

Glass, A. L., Holyoak, K. J., & O'Dell, C. Production frequency and the verification of quantified statements. *Journal of Verbal Learning and Verbal* Behavior, 1974, **13,** 237–254.

Guzmán, A. Decomposition of a visual scene into three-dimensional bodies. In A. Grasselli (Ed.), *Automatic interpretation and classification of images.* New York: Academic Press, 1969.

Hubel, D. H., & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology (London),* 1962, **160,** 106–154.

Hubel, D. H., & Wiesel, T. N. Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology (London),* 1968, **195,** 215–243.

Jacoby, L. L. Encoding processes, rehearsal and recall requirements. *Journal of Verbal Learning and Verbal Behavior,* 1973, **12,** 302–310.

Jacoby, L. L. The role of mental contiguity in memory: Registration and retrieval effects. *Journal of Verbal Learning and Verbal Behavior,* 1974, **13,** 483–496.

Johnson, N. F. The role of chunking and organization in the process of recall. In G. H. Bower (Ed.), *The psychology of learning and motivation.* Vol. 4. New York: Academic Press, 1970.

Kahneman, D., & Tversky, A. Subjective probability: A judgment of representativeness. *Cognitive Psychology,* 1972, **3,** 430–454.

Kahneman, D., & Tversky, A. On the psychology of prediction. *Psychological Review,* 1973, **80,** 237–251.

Kintsch, W. *The representation of meaning in memory.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1974.

Lashley, K. S. The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior.* New York: Wiley, 1951.

Lesgold, A. N. & Bower, G. H. Inefficiency of serial knowledge for associative responding. *Journal of Verbal Learning and Verbal Behavior,* 1970, **9,** 456–466.

Lindsay, P. H., & Norman D. A. *Human information processing: An introduction to psychology.* New York: Academic Press, 1972.

Locke, E. A. Is "behavior therapy" behavioristic? *Psychological Bulletin,* 1971, **76,** 318–327.

Mandler, G. Organization and memory. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation.* Vol. 1. New York: Academic Press, 1967.

Maslow, A. H. A theory of metamotivation: The biological rooting of the value life. *Journal of Humanistic Psychology,* 1967, **7,** 93–127.

Meacham, J. A. The development of memory abilities in the individual and in society. *Human Development,* 1972, **15,** 205–228.

Meyer, B. J. F. The organization of prose and its effect on recall. Technical Report No. 1 (Ph.D. dissertation) (Reading and Learning Series), Department of Education, Cornell University: Ithaca, New York, 1974.

Miller, G. A. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review,* 1956, **63,** 81–97.

Miller, G. A., Galanter, E., & Pribram, K. H. *Plans and the structure of behavior.* New York: Holt, Rinehart & Winston, 1960.

Miller, N. E. Some reflections on the law of effect produce a new alternative to drive reduction. In M. R. Jones (Ed.), *Nebraska symposium on motivation.* Lincoln: Nebraska University Press, 1963. Pp. 65–112.

Mischel, W. Continuity and change in personality. *American Psychologist,* 1969, **24,** 1012–1018.

Mischel, W. Towards a cognitive social learning reconceptualization of personality. *Psychological Review,* 1973, **80,** 252–283.

Morton, J. The interaction of information in word recognition. *Psychological Review,* 1969, **76,** 165–178.

Newell, A., *et al. Speech understanding systems: Final report of the study group.* New York: American Elsevier, 1973.

Newell, A., & Simon, H. A. *Human problem solving.* Englewood Cliffs, New Jersey: Prentice-Hall, 1972.

Norman, D. A., & Rumelhart, D. E. *Explorations in cognition.* San Francisco: Freeman, 1975.

Parkinson, S. R., Parks, T. E. & Kroll, N. E. A. Visual and auditory short-term memory. *Journal of Experimental Psychology,* 1971, **87,** 274–280.

Pew, R. W. Human perceptual–motor performance. Technical Report No. 48. Human Performance Center, University of Michigan, Ann Arbor, Michigan, August, 1974.

Postman, L. Transfer, interference and forgetting. In J. W. Kling & L. A. Riggs (Eds.), *Woodworth and Schlosberg's experimental psychology.* (3rd ed.) New York: Holt, Rinehart & Winston, 1971.

Prytulak, L. S. Natural language mediation. *Cognitive Psychology,* 1971, **2,** 1–56.

Ratliff, F. *Mach bands: Quantitative studies on neural networks in the retina.* San Francisco: Holden-Day, 1965.

Reed, S. K. Pattern recognition and categorization. *Cognitive Psychology,* 1972, **3,** 382–407.

Rescorla, R. A. Informational variables in Pavlovian conditioning. In G. H. Bower (Ed.), *The psychology of learning and motivation.* Vol. 6. New York: Academic Press, 1972. Pp. 1–46.

Rieger, C. J. Conceptual memory: A theory and computer program for processing the meaningful content of natural language utterances. Unpublished doctoral dissertation, Stanford University, 1974.

Rumelhart, D. E. Notes on a schema for stories. Unpublished working paper, Psychology Department, University of California at San Diego, October, 1973.

Rundus, D. Analysis of rehearsal processes in free recall. *Journal of Experimental Psychology,* 1971, **89,** 63–77.

Schank, R. C. Conceptual dependency: a theory of natural language understanding. *Cognitive Psychology*, 1972, **3**, 552–631.

Schank, R. C. The development of conceptual structures in children. Technical Report AIM-203, Artificial Intelligence Laboratories, Computer Science Department, Stanford University, May 1973. (a)

Schank, R. C. Identification of conceptualizations in natural language. In R. C. Schank & K. M. Colby (Eds.), *Computer models of thought and language.* San Francisco: Freeman, 1973. (b)

Shanteau, J. Component processes in risky decision making. *Journal of Experimental Psychology*, 1974, **103**, 680–691.

Shiffrin, R. N., & Atkinson, R. C. Storage and retrieval processes in long-term memory. *Pshchological Review*, 1969, **76**, 179–193.

Simon, H. A. Motivational and emotional controls of cognition. *Psychological Review*, 1967, **74**, 29–39.

Skinner, B. F. *About behaviorism.* New York: Knopf, 1974.

Slovic, P., & Lichtenstein, S. Relative importance of probabilities and payoffs in risk taking. *Journal of Experimental Psychology Monograph*, **78**, (No. 3, Part 2), 1968.

Sternberg, S. The discovery of processing stages: extensions of Donder's method. Attention and Performance. *Acta Psychologica*, 1969, **30**, 276–315.

Taylor, C. *The Explanation of Behavior.* London: Routledge & Kegan, Paul, 1964.

Thorndike, E. L. *Human learning.* New York: Century Co., 1931. (Paperback ed., Cambridge, Massachusetts: MIT Press, 1966)

Tolman, E. C. *Purposive behavior in animals and men.* Berkeley: University of Calfornia Press, 1932.

Tulving, E. Subjective organization and effects of repetition in multi-trial free recall learning. *Journal of Verbal Learning and Verbal Behavior*, 1966, **5**, 193–197.

Turvey, N. T. Peripherial and central processes in vision: Inferences from an information processing analysis of masking with pattern stimuli. *Psychological Review.* 1973, **80**, 1–52.

Tversky, A., & Kahneman, D. Belief in the law of small numbers. *Psychological Bulletin*, 1971, **76**, 105–110.

Tversky, A., & Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 1973, **5**, 207–232.

Tversky, B. Encoding processes in recognition and recall. *Cognitive Psychology*, 1973, **5**, 275–324.

Vicens, P. Aspects of speech recognition by computer. Technical Report CS-127, doctoral thesis, Computer Science Department, Stanford Univ., 1969.

von Békésy, G. *Sensory inhibition.* Princeton, New Jersey: Princeton University Press, 1967.

von Restorff, H. Analyse von Vorgangen in Spurenfeld. I. Über die Wirkung con Bereichsbildungen im Spruenfeld. *Psychologisch Forschung*, 1933, **18**, 299–342.

Waugh, N. C., & Norman D. A. Primary memory. *Psychology Review*, 1965, **72**, 89–104.

Wescourt, K. T., & Atkinson, R. C. Scanning for information in long- and short-term memory. *Journal of Experimental Psychology*, 1974, **98**, 95–101.

Winston, P. H. Learning structural descriptions from examples. MIT Artificial Intelligence Laboratory. Project AI-TR-231, 1970.

3

Cognitive Theory Applied to Individual Differences

Earl Hunt

Marcy Lansman

The University of Washington

All men may be created equal, but no two are created identical. Over a hundred years ago Sir Francis Galton accepted this fact as a central problem for psychology. He devoted much of his life to an attempt to explain individual differences in mental ability in terms of general psychological laws. For a number of historical and social reasons Galton's goal of scientific explanation was put aside in order to develop a sophisticated system for measuring cognitive power without attempting to explain it (Tyler, in press). The result is the intelligence test that we know today, and which has become our normal way of describing differences in cognition. Yet the questions Galton posed are important scientific issues. We shall try to restate them in terms of the modern theory of cognitive psychology.

What might constitute a theory of individual differences? Since the early 1900s, great emphasis has been placed upon classification of the basic ways in which individuals' thought processes differ. Classification is certainly a legitimate stage in the development of a scientific theory. However, we part company with the traditional theory of mental testing, or *psychometrics,* in stating that we do not regard a classification scheme as an adequate final goal. Instead we want a theory which describes the *processes* by which different people attack the same problem.

To illustrate what is meant, consider an analogy to the measurement of automobile abilities. An observer could develop a classificatory scheme