

COMPUTER SIMULATION OF HUMAN THINKING
AND PROBLEM SOLVING

Allen Newell
Computer Sciences Department
The RAND Corporation

H. A. Simon*
Consultant, The RAND Corporation

P-2312

May 29, 1961

*Professor, Carnegie Institute of Technology,
Pittsburgh, Pennsylvania

This paper was based on a talk given on March
23, 1961 by H.A. Simon in the Massachusetts
Institute of Technology Centennial Lecture
Series on "Management and the Computer of the
Future."

Reproduced by

The RAND Corporation • Santa Monica • California

The views expressed in this paper are not necessarily those of the Corporation

COMPUTER SIMULATION OF HUMAN THINKING AND PROBLEM SOLVING

INTRODUCTION

The use of computers to simulate human thinking has a prehistory and a history, as well as a present and a future. During the prehistorical period, prior to World War II, there were no computers in the modern sense, but there were a number of successful attempts to construct teleological mechanisms--analogue devices that simulated one aspect or another of an organism's adaptive behavior in relation to its environment.¹

History begins in earnest, however, with the rapid growth of servo-mechanism theory during World War II and with the appearance of the first stored-program digital computers--two of the three legs on which Professor Wiener's cybernetics stands. Grey Walter's "tortoises" and W. Ross Ashby's Homeostat represent important early progress, as does an analogue simulation of a self-organizing network that Professor Minsky constructed in 1951.

Computer simulation had already begun to take definite form as a field of research by the time of the well-known session on learning machines at the 1955 Western Joint Computer Conference, a session in which Professor Miller also

¹A number of these undertakings are catalogued in Professor Boring's instructive and entertaining paper, "Mind and Mechanism," Journal of Psychology, 59 (1946), 173-192.

participated.² At that session, Clark and Farley of Lincoln Laboratories described a computer simulation of a self-organizing "nerve net" system; Selfridge and Dinneen, also of Lincoln Laboratories, described a pattern recognition program; and one of the authors, Newell, of The RAND Corporation, outlined a program for a chess-playing machine.

One of the discussants at that session, Walter Pitts, observed that there were two main lines of attack represented: the first taking as its point of departure some features of the human nervous system and sensory apparatus, the second, the organization of symbolic processes to perform complex thinking tasks. As Mr. Pitts put it:³

The speakers this morning are all imitators in the sense that the poet in Aristotle "imitates" life. But, whereas Messrs. Farley, Clark, Selfridge, and Dinneen are imitating the nervous system, Mr. Newell prefers to imitate the hierarchy of final causes traditionally called the mind. It will come to the same thing in the end, no doubt ...

Most workers in this field continue to believe that it will come to the same thing, but the end is not yet, and these two main strands of research are still clearly discernable in the work going on at the present time. Our remarks this

²See "Session on Learning Machines," Proceedings of the 1955 Western Joint Computer Conference, Institute of Radio Engineers, 1955, pp. 85-111.

³Ibid., p. 108.

evening will be concerned almost exclusively with the second strand--with the imitation of mind. This strand has already begun to make contact with important potential areas of application, business administration and teaching among them.⁴ Our purpose tonight, however, is not to speculate about applications. We shall be speculative enough, we are sure, for your tastes, but we shall speculate about the form that fundamental theory in this field is taking, rather than about the implications of that fundamental theory for everyday affairs.

THE PROOF OF POSSIBILITY

It is no longer necessary to argue that computers can be used to simulate human thinking, or to explain in general terms how such simulation can be carried out. A dozen or more computer programs have been written and tested that perform some of the interesting symbol-manipulating, problem-solving tasks that humans can perform, and do so in a manner which simulates, in some general respects, the way humans do these tasks. Computer programs now play chess and checkers, find proofs for theorems in geometry and logic, compose music, balance assembly lines, design electric motors and generators, memorize nonsense syllables, form concepts,

⁴See H. A. Simon, The New Science of Management Decision (New York: Harper and Brothers, 1960).

and learn to read.⁵

With the proof of possibility accomplished, we can turn to more substantive questions. We can ask what we have learned about human thinking and problem solving through computer simulation: to what extent we now have theories for these phenomena, and what the content of these theories is. Since we want to talk about these substantive matters, we shall simply make the following assertions, which are validated by existing computer programs.

1. Computers are quite general symbol-manipulating devices that can be programmed to perform nonnumerical as well as numerical symbol manipulation.

2. Computer programs can be written that use nonnumerical symbol manipulating processes to perform tasks which, in humans, require thinking and learning.

3. These programs can be regarded as theories, in a completely literal sense, of the corresponding human processes. These theories are testable in a number of ways: among them, by comparing the symbolic behavior of a computer so programmed with the symbolic behavior of a human subject when both are performing the same problem-solving or thinking tasks.

THE GENERAL PROBLEM SOLVER

The theory we shall have most to say about is a computer program called the General Problem Solver. It is not "general"

⁵For an excellent recent survey of heuristic programs, although with emphasis upon "artificial intelligence" rather than simulation of human thought, see Marvin Minsky, "Steps Toward Artificial Intelligence," Proceedings of the Institute of Radio Engineers, 49 (January 1961), 8-30.

in the sense that it will solve, or even try to solve, all problems--it obviously won't. It is called "general" because it will accept as tasks all problems that can be put in a specified, but fairly general, form, and because the methods it employs make no specific reference to the subject matter of the particular problem it is solving. The General Problem Solver is a system of methods--believed to be those commonly possessed by intelligent college students--that turn out to be helpful in many situations where a person confronts problems for which he does not possess special methods of attack.

Before general methods can be applied to any particular class of problems, of course, the problem solver must also learn, or be taught, the rules that apply to that particular problem domain. The General Problem Solver will not prove theorems unless instructed in the rules of proof in the particular branch of mathematics to which the theorems belong. Thus, in any particular problem domain, the resources available to the General Problem Solver include information about the task environment as well as its own repertory of methods.

Missionaries and Cannibals

Let us introduce the General Problem Solver (which we shall call GPS) by means of a simple example. Many of you are familiar with the puzzle of the Missionaries and Cannibals, and some of you saw a young lady solving the puzzle in a recent CBS television program celebrating MIT's centenary. There are

three missionaries and three cannibals on the bank of a wide river, wanting to cross. There is a boat on the bank, which will hold no more than two persons, and all six members of the party know how to paddle it. The only real difficulty is that the cannibals are partial to a diet of missionaries. If, even for a moment, one or more missionaries are left alone with a larger number of cannibals, the missionaries will be eaten. The problem is to find a sequence of boat trips that will get the entire party safely across the river--without the loss of any missionaries.

Suppose, now, that we encountered this puzzle for the first time. We are endowed by nature and nurture with certain abilities that enable us to tackle the problem. We might or might not solve it, but we could at least think about it. In what would this thinking consist? In particular, how could we bring to bear our general problem-solving skills, which make no reference to missionaries and cannibals, on this particular situation?

Clearly, we have to form some kind of abstraction of the problem that will match the abstractness of our general methods: We have some people and a boat on this side of the river, and we want them on that side of the river. Stated abstractly, we have a certain state of affairs, and we want a different state of affairs. Moreover, we can describe both states and we can also describe what the differences are between them--between what we have and what we want.

In this case, the differences between the given and the desired are differences in physical location. Our men are on one side of the river; we want them on the other. But we have had vast experience with differences in location, and that experience (stored somehow in memory) tells us that boats are useful devices for reducing differences of location on water. So we begin to consider the possible sequences of boatloads that will get our party across the river without casualties.

It is clear from this formulation of the problem what part is played in its solution by our general problem-solving techniques and what part by our knowledge and experience of the particular problem domain in question. A general solution technique is to characterize the given and desired situations, to find the differences between them, and to search for means--implements or operators--that are relevant to removing differences of these kinds. Our knowledge of the task and our experience tell us what the given and desired situations are, and what kinds of operators may be relevant for getting from here to there.

Structure of GPS

We can now characterize the program of the General Problem Solver more formally.⁶ The program deals with symbolic objects

⁶For a fuller description, see A. Newell, J. C. Shaw, and H. A. Simon, "Report on a General Problem-Solving Program," in Information Processing, Proceedings of the International Conference on Information Processing, (Paris, UNESCO, 1960) 256-264.

that describe or characterize situations--the given situation, the desired situation, or various intermediate possible situations. The program also deals with symbols representing differences between pairs of objects, and with symbols representing operators that are capable of inducing changes in the objects to which they are applied. (See Fig. 1, left-hand column).

Goal Types. The processes of GPS are organized around goals of three types:

1. Transformation goals: to transform object a into object b.
2. Difference reduction goals: to eliminate or reduce difference d between objects a and b.
3. Operator Application goals: to apply operator q to object a.

Methods. With each type of goal in GPS there is associated one or more methods, or processes, that may contribute to the attainment of the goal. The principal methods in the present version of GPS are three in number, one for each goal type:

1. Method for transformation goals: to transform a into b,
 - a. Notice a difference, d, between a and b;
 - b. Establish the goal of reducing d between a and b;
 - c. Try to attain this new goal;
 - d. If successful, find a new difference and repeat.
2. Method for difference reduction goals: to reduce d between a and b,

GPS	LEARNING SPEECH	ADAPTING TO ENVIRONMENT
Objects	Perceptual symbols (audited phonemes)	AFFERENT
Differences	Comparison between adult-child phoneme images	State language
Relevant Operators	Changes in motor symbols (control of speech production)	EFFERENT Process language

Fig. 1

COMPARISON OF BASIC CATEGORIES IN
GPS, SPEECH LEARNING,
AND ORGANISMIC ADAPTATION

- a. Recall an operator, q, that is relevant to differences of the type of d;
 - b. Establish the goal of applying q to a;
 - c. Try to attain this new goal;
 - d. If successful, return to the previous transform goal.
3. Method for operator application goals: to apply operator q to a,
- a. Compare conditions for application of q with object a;
 - b. If these are not satisfied, establish and try to attain the goal of transforming a into an object that meets these conditions;
 - c. When the conditions are satisfied, apply q to a, and return to the previous difference reduction goal with the modified object, a'.

This is a rather simplified description of what goes on in GPS, but it gives the broad outline of the program. GPS, to put it simply, is a program that reasons about ends and means. It is capable of defining ends, seeking means to attain them, and, in the process of so doing, defining new subsidiary ends, or subgoals, towards the original end.

As a theory of human problem solving, GPS asserts that college students solve problems--at least problems of the sorts for which the program has been tested--by carrying out this kind of organized ends-means analysis. It does not

assert that the process is carried out consciously--it is easy to show that many steps in the problem-solving process do not reach conscious awareness. Nor does the theory assert that the process will appear particularly orderly to an observer who does not know the program detail or, for that matter, to the problem solver himself. It does assert that if we compare that part of the human subject's problem solving behavior which we can observe--the steps he takes, his verbalizations--with the processes carried out by the computer, they will be substantially the same.

Abstracting and Planning Processes. Before we leave this description of GPS, we should like to mention one other kind of process that we are incorporating in the program, and that certainly must be included if we are to explain and predict the behavior of our subjects--particularly the brighter ones. We call these additional methods abstracting and planning processes. Briefly, abstracting consists in replacing the objects, the differences, and the operators, with new symbolic expressions that describe the situation in much more general terms, omitting the detail.⁷ For example, we might ask GPS to prove a trigonometric identity:

$$\cos^2 x + \sin^2 x = \tan x \cot x.$$

⁷See ibid., pp. 261-2, for a description of a specific planning method for GPS. In our subjects, abstracting often takes the form of simply ignoring some of the problem detail at certain stages of the solution process.

Here, GPS might take as a the expression, " $\cos^2x + \sin^2x$," and as b the expression, " $\tan x \cot x$." In using the planning method, these might be abstracted to: (a') "an expression containing cos and sin" and (b') "an expression containing tan and cot," respectively. Then, the methods of GPS could be applied to transforming the abstracted given object, a', into the abstracted desired object, b'. If this goal were attained, the steps employed for this transformation would generally provide a plan for transforming the original, detailed given object, a, into the original desired object, b. In the particular case illustrated, the plan might be something like: "First eliminate cos and sin from the expression, and then introduce tan and cot."

The Generality of Ends-Means Analysis

The processes incorporated in GPS have actually been observed in the behavior of our human subjects solving problems in the laboratory. By analysing the tape-recorded protocols of their problem-solving efforts, we can identify the occurrences of the three goal types and the four methods. Moreover, the augmented GPS, containing the planning method, incorporates a substantially adequate set of processes to explain our subjects' behavior in some of these simple theorem-proving, puzzle-solving situations.⁸ By the adequacy of GPS, we mean two things:

⁸See A. Newell and H. A. Simon, "The Simulation of Human Thought," in Current Trends in Psychological Theory (Pittsburgh: U. of Pittsburgh Press, 1961).

1. We do not find in the subjects' protocols evidences of processes quite different from those postulated in GPS. This may mean only that we don't know how to look for them, but,

2. When we have compared the trace of the GPS computer (or hand simulations of the computer program) with the protocols of a subject solving the same problem, we have found that the two often follow the same path--noticing the same things about the problem expressions, establishing the same subgoals, applying the same operators, running down the same blind alleys--over periods of time ranging up to several minutes. That is to say, the processes in GPS are sufficient to produce a stream of behavior in a given problem situation quite similar to that produced by the human subject.

These kinds of tests, even if broadened, would still not say much about the generality of GPS as a theory of human thinking and problem solving. It might turn out that if we examined tasks quite different from those used in developing the program, and made the same careful records of subjects' protocols, we would find many new processes exhibited that are not contained in GPS. However, extensions of GPS in fair detail to problem domains that were not considered when the program was developed indicate that its processes are adequate at least to these other domains. For example, Missionaries and Cannibals, which was first suggested as a possible task by Mr. Thomas Wolf of the Columbia Broadcasting System, has

been solved by the current version of GPS--not without some reorganization of the program, but without addition of new goal types or methods. Similarly, the applications to algebraic and trigonometric identities and to certain learning tasks appear to require no enlargement of the basic repertory of methods. Less detailed analysis of a variety of other tasks shows GPS to be adequate for these also.

Still, these additional tests do not carry GPS beyond a fairly limited range of formal problem-solving situations. It would be of considerable interest to explore, even qualitatively, the powers and limitations of GPS when it is confronted with a thinking or learning task of quite a different kind from any of these. We should like, this evening, to carry out a reconnaissance along these lines. First, we will describe, on the basis of what is now known, the processes that humans use in a task that appears, superficially, to be quite different from problem solving. Then, we shall propose a framework which shows that these processes can be subsumed under those already incorporated in the General Problem Solver. The particular task we shall examine was chosen because quite a bit is known about it, and because it will allow us to call on our discussants for a maximum of assistance this evening.

THE ACQUISITION OF SPEECH

There are many human activities to which we would apply the term "thinking" but not the term "problem solving." There

are also many activities we would usually call "learning" rather than "thinking." We would ordinarily call a child's acquisition of speech, "learning." We propose to consider the acquisition of speech as an example of human cognitive activity that is almost at an opposite pole from the rather highly verbalized, somewhat conscious, practiced problem solving of an intelligent and educated adult. We can then judge whether the processes at these two poles are quite different or basically the same.

Speech acquisition has been about as well studied as any non-laboratory complex human activity, and a review of the literature indicates that there is general consensus about the particular facts we shall use.⁹

Central Representations

We consider an infant who has already learned the names of a few objects--as evidenced by the fact that he can point to them or fetch them when they are named by an adult--but who has not yet pronounced their names. From his behavior, we can infer that when the child perceives the spoken word "ball," his perception has some kind of internal representation in the

⁹See, for example, C. E. Osgood, Method and Theory in Experimental Psychology, (New York: Oxford University Press, 1953) 683-690; G. A. Miller, "Speech and Language," Chapter 21 in S. S. Stevens, ed., Handbook of Experimental Psychology (New York: Wiley, 1951); and G. A. Miller, Language and Communication (New York: McGraw-Hill, 1951), Chapter 7.

brain that permits it to be associated, through previous experience, with some internal representation of a visually perceived ball.

To say the word "ball," the child must, in addition, store some kind of program capable of energizing, through motor (efferent) channels, the muscles involved in speech production-- in the production of the specific phonemes of that word. Let us call the "whatever-it-is" in the central nervous system that represents internally a perceived sensory stimulus an afferent or perceptual symbol. Let us call the "whatever-it-is" that represents the program for initiating the motor signals an efferent or motor symbol.

Learning to speak, in this formulation, means acquiring the motor symbols that correspond to perceptual (auditory) symbols of words already known, and associating the former with the latter. Now the difficulty is that there is no way in which the corresponding perceptual and motor symbols can "resemble" each other--can symbolize the appropriateness of their association by resemblance. The correspondence is purely arbitrary.¹⁰ The infant is faced (if he only knew it!) with the immense inductive task of discovering which motor symbols will cause speech production that, when he hears it, will produce, in turn, an appropriate auditory symbol to be per-

¹⁰We shall have occasion to qualify the adverb "purely" when we come to consider the factorization of words into phonemes and phoneme components.

ceived and recognized. And the task appears at first blush to have little structure that would permit it to be approached with some less arduous technique than trial-and-error search.

There is ample evidence that much trial-and-error search is indeed required before the infant acquires the skill of speaking. The child imitates the adults around him, and he imitates himself (echoic speech). Gradually, over many months, he acquires the motor symbols that enable him to produce sounds which he hears as the expected auditory symbols. In the early stages, the child's acquisition of a speaking vocabulary appears to be paced by the task of developing the new motor symbols. At later stages, he is able to produce a word relatively easily once he has learned to recognize the corresponding auditory symbol.

Factorization

A little reflection will persuade us that something more than trial and error is involved. If that were all, the three hundredth word would be no easier to pronounce than the first. The child learns to learn. In what does this consist?

Although the motor symbol cannot be compared with the perceptual symbol, the correct perceptual symbol for a word can be compared, through imitation, with the perceptual symbol produced by the attempt to pronounce the word. If these are different, modification of the motor symbol can be attempted until an auditory symbol resembling the correct one is perceived.

Thus far we have been assuming that the units in terms of which these transactions take place are words. But there is no reason for this assumption--the child might well attend to particular syllables, phonemes, or even components of phonemes. The auditory symbols for words can be compound symbols or expressions--strings of phonemes, each phoneme itself encoded in terms of its component frequencies and other characteristics. It is even more plausible to suppose that the motor symbols would be constructed from smaller units, for each word involves a temporal succession of syllables, each syllable a temporal succession of phonemes, and each phoneme a whole set of signals to the several muscles involved in that part of the speech act. Thus, one of the many components of the motor symbol for the spoken word "dog" might be the signal that pushes the tongue against the palate in the initial "d" phoneme of this one-syllable word.

The Learning Process

There is considerable evidence today that this picture of the processes of word-recognition and word-production is correct, at least in broad outline. Many of the components involved in both auditory and motor symbols have been tentatively identified, and there is good experimental evidence for some

of them.¹¹ But what does the picture, if true, contribute to our understanding of the child's acquisition of speech?

It means that the inductive learning need not be blind inductive learning--attempting to associate by pure trial and error each of a large number of words with an appropriate motor symbol chosen from the myriad of producible sequences of speech sounds. On the contrary, to the extent that specific factors in the auditory symbol vary with specific factors in the motor symbol (e.g., as one of the formant frequencies in vowel sounds varies with the size of the resonating mouth cavity), the search for the correct symbol can be very much restricted. Components can be corrected on a one-at-a-time basis. For example, the child trying to pronounce "dog" can at one time attend to the correctness of the vowel, at another time to the correctness of the initial consonant, or even to the aspect of the initial consonant associated with tongue position.

Thus, the hypothesis of factorization is supported both by experimental evidence that it does take place, and by theore-

¹¹For a general introduction to these topics, see G. A. Miller, Language and Communication, (New York: McGraw-Hill, 1951) Chapter 2. An excellent recent survey is Richard Fatchchand, "Machine Recognition of Spoken Words," Advances in Computers, F. L. Alt, ed., (New York: Academic Press, 1960), 193-321. See also J. W. and C. D. Forgie, "Results Obtained from a Vowel Recognition Computer Program," The Journal of the Acoustical Society of America, 31 (1959), 1480-89, and A. M. Liberman, et al., "Minimal Rules for Synthesizing Speech," ibid., 31 (1959), 1490-99. The last three references cited illustrate, incidentally, the large role that computers are playing today in linguistic and phonetic research.

tical reasons why it "should" take place--why speech acquisition would be very much easier with it than without it. Trial-and-error acquisition of words without factorization would require a search, in each instance, for the correct motor symbol from among tens of thousands of possible symbols. Trial-and-error acquisition of phonemes would require a search from among only a few hundred phonemes (much fewer are actually used, of course, in any single dialect). Trial-and-error search among phoneme components would be even more restricted--there are, for example, probably only a half dozen distinguishable tongue positions. Thus, by factorization of the total space of possibilities, a very limited trial-and-error search of the factors can be substituted for an immense search of the product space. Moreover, once the child has acquired motor symbols corresponding to the common phonemes, acquisition of new words (new combinations of these same phonemes) can be very rapid.

Summary: The Child's Acquisition of Speech

Let us now summarize our description, partly factual, partly hypothetical, of the speech acquisition process. The child acquires perceptual auditory symbols corresponding to words he has heard and has associated with visual symbols. He tries, on a trial-and-error basis, to produce words, hears his productions, and compares these auditory symbols with those already stored. When he detects differences, he varies the motor symbol to try to remove them. As he learns, he detects that changes

in certain components of the motor symbols alter only certain components of the auditory symbols. Thus, he is able to factor the correction process and, thereby, accelerate it greatly.

Acquisition of Speech by GPS

Now it is very easy, with a few changes in vocabulary, to translate this whole description back in terms of GPS. When the translation has been made, we shall see that the processes just described are the methods of GPS.

Let us, in this translation, call the auditory symbols objects. (Figure 1, second column.) We assume that there exist central processes that modify motor symbols--that change one or more of their components. We will call these processes operators. A change in a motor symbol will, in turn, change the auditory symbol that is perceived when that motor symbol produces a sound.

The child detects differences between the object he has produced (i.e., his perception of the sound) and the correct object (his perception of the sound when produced by adults). He applies operators to the motor symbol to modify the sounds he produces, hence the object perceived, and he compares the latter again with the correct object. This search process continues until he can reproduce the perceived object.

But this does not account for the factorization, which we have argued is so crucial to the efficiency of the learning process. How will GPS learn (1) which differences in objects

are associated with which operators upon the motor symbols, and (2) how to factor objects and operators? Although the answers to these questions are far from certain, a scheme we have proposed elsewhere would enable GPS to handle these tasks also.¹² We will sketch it briefly:

1. Given a set of differences and a set of operators, GPS can, with modest amounts of trial and error, detect which operators are relevant to producing or eliminating which differences. To take a crude, but simple, example: it takes relatively little trial and error to discover what differences in the perceived sound are associated with changes in the rounding of the lips while producing a vowel. The factorization has already largely been carried out by nature, so to speak, because changes in only a few aspects of the motor signal will change only a few aspects of the perceptual symbol.

2. The GPS processes can themselves be employed to discover inductively a "good" factorization--a "good" set of differences. To do this, GPS must be supplied with some very general criteria as to what constitutes such a good set. The criteria would be of the following general kind:

¹²The full account of this learning scheme is given in A. Newell, J. C. Shaw, and H. A. Simon, "A Variety of Intelligent Behavior in a General Problem Solver," pp. 153-187 in M. C. Yovits and S. Cameron, eds., Self-Organizing Systems, (New York: Pergamon Press, 1960).

- a. Only one or a few operators should be relevant to each difference (so that, given a difference, an appropriate operator can be found without too much search).
- b. Only one or a few differences should be associated with each operator (so that the sounds produced can be varied factor by factor).

With such a set of criteria provided, finding a good set of differences simply becomes another kind of problem to which GPS can apply its problem-solving methods. What are the objects, differences, and operators in terms of which this new kind of problem is formulated? To avoid unnecessary confusion, we will capitalize the terms OBJECTS, DIFFERENCES, and OPERATORS in speaking of the new problem context in order to distinguish them from the objects (perceptual symbols), differences, and operators (changes in motor symbols) involved in the original task of acquiring speech.

The OBJECTS for the new problem-solving task are the sets of differences in the original task environment. The new DIFFERENCES designate to what extent particular sets of differences meet the criteria we have just listed. OPERATORS are processes for altering the set of differences under consideration by deleting differences from the set, adding differences, or generating new differences for possible inclusion. GPS then tests in what respects particular OBJECTS (sets of differences)

are DIFFERENT from the desired OBJECT (as indicated by the criteria). It seeks to remove these DIFFERENCES (modify the set of differences) by applying OPERATORS (by adding, subtracting, or modifying differences).

Since this scheme has not been realized on a computer, we cannot tell how effective GPS would be in handling it. All we can say is that it is a problem whose solution can be attempted with the means at the disposal of GPS.

A due respect for parsimony would suggest, then, that instead of postulating quite different processes for the acquisition of such skills as speaking from those postulated for adult problem solving, we embrace tentatively the hypothesis that the processes are in fact the same--that the General Problem Solver provides a description of both processes. This hypothesis would provide a sharp focus for empirical research into the early speech behavior of the child.

THE STATE-PROCESS DICHOTOMY

Let us accept this hypothesis for the moment; that the same system of ends-means processes is involved in learning speech and in problem solving. Can we explain why a system of ends-means analysis should provide the basis for adaptive behavior in both classes of situations? We shall try to provide an explanation for the generality of ends-means processes by showing how these arise quite naturally from the problem that any organism must solve if it is to use its sensory and motor apparatus effectively to survive.

Relation of Perceptual to Motor Symbols

The terms "perceptual" and "motor," or "afferent" and "efferent," reflect the dual relation that every adaptive organism has with its environment. It perceives aspects of the environment, and it acts upon the environment. It must be able, therefore, to transmit, store, and operate upon internal representations (perceptual symbols) that stand for its perceptions, and it must be able to transmit, store, and operate upon internal representations (efferent or motor symbols) that can serve as signals to its effectors. The organism survives by associating appropriate motor symbols with the perceptual symbols that stand for various classes of perceptions.¹³

In particular, the organism can perceive, at least grossly, its own behavior caused by its efferent signals. Hence, among the perceptual symbols that it can store are symbols that stand for the perception of corresponding motor signals. Languages are especially adapted to facilitate this correspondence. Language behavior, built from limited alphabets of unit behaviors, is highly stylized, so that to each distinct language "act" will correspond an easily perceivable and distinguishable perceptual symbol.

¹³We need hardly say that this description does not commit us to any over-simplified reflex-arc picture of the peripheral and central systems. GPS is a concrete example of a system of the sort we are describing. In it, the perceptual-motor associations are represented by the table of connections between differences and operators. The use it makes of these connections, and, consequently, the relation of response to stimulus is highly complex.

Nevertheless, the relation of a particular language efferent--say that which energizes the word "dog"--to the corresponding perceptual symbol is arbitrary. There is no more resemblance between the auditory "dog" and the motor symbol which produces that word than between "dog" and "Hund." If it is to be learned, the correspondence must be learned as a pure fact. By building up a dictionary relating motor with perceptual symbols--including language-symbols--the organism gains the ability to produce the actions it "intends." In the last section we explored how this ability could develop in the case of speech.

The duality of our relation with the environment reveals itself in the vocabulary of natural language--particularly in the distinction between nouns and adjectives, on the one hand, and verbs, on the other. We have clean clothes (a perceptual symbol) because they have been washed (a motor symbol). It is a fact stored in our "table of connections" that when we wash clothes they become clean. As we build up our vocabulary, however, we pass more readily from the one mode of discourse to the other. Thus, the clothes, in the last example, might also have been cleaned. As we learn what actions have what effects, changes in objects are named by the processes that produce them, and processes by the effects they create.

The Problem of Translation

It is precisely this duality of language--or, more broadly, of the internal symbols employed in thought--that makes behavior problematic. The world as it is and as it is desired is described in a state language, a language of perceptual symbols.¹⁴ Possible actions are described in a process language, a language of motor symbols. (Fig. 1, third column.) The problem of adapting is the problem of finding the statement in the process language that corresponds to the difference between existing and desired states of affairs in the state language.¹⁵

But the problems that GPS was designed to handle can be viewed in exactly the same way. What is involved in discovering a proof for the Pythagorean Theorem? The theorem is a symbolic object in the state language: "The square on the hypotenuse of a right triangle is equal to the sum of the squares on the sides." By comparing this theorem, so stated, with the axioms and previously proved theorems, we detect differences between them. A proof of the theorem is a symbolic object in the process language. This object--the justification that

¹⁴It should be observed that the body is part of the environment that is perceived. Hence, drives like hunger produce perceptual symbols just as external senses do. Or perhaps it would be better to say that the drive is the perceptual symbol produced by the perception of hunger.

¹⁵Other aspects of the state-process dichotomy are discussed in A. Newell, J. C. Shaw, and H. A. Simon, "The Processes of Creative Thinking," presented before a symposium on creativity at the University of Colorado, May 14, 1958, (The RAND Corporation, Paper P-1320, September 16, 1958).

we generally write down alongside the successively modified axioms and theorems--describes the sequence of operations that eliminates the differences between axioms and desired theorems. Given a set of axioms, for every theorem defined in the state language, the theorem can be represented in the process language by the sequence of operations that constitutes its proof.

Thus mathematics, and problem-solving generally, is an imitation of life. Problem-solving activity uses the very fundamental processes that all adaptive organisms must have if they are to coordinate successfully their perceptual and motor pictures of the world.¹⁶ Ends-means relations, far from being highly special, are reflections of the basic state-process dichotomy, the dichotomy between perceiving and acting.

The Difficulty of the Environment

How hard a problem will be depends on the simplicity or complexity of the rules that define the correspondence between the two languages. An example of a relatively simple

¹⁶We should like to call attention to the similarity, hardly accidental, between the translation problems we have been describing and some of the most striking results of modern mathematics--the undecidability theorems. Decision problems, in the Gödel sense, can always be represented as problems of finding in language B the representation of an object that is described in language A, where there are at least some objects that have names in both languages. If the rules of correspondence between the two languages are sufficiently complicated, this will be difficult, and, as the undecidability theorems show, it may be impossible.

correspondence is the relation between the decimal and octal representations of integers. There is a simple and direct algorithm that solves all problems of the form: if a is the decimal representation of a number, what is its octal representation?

On the other hand, the correspondence between the vocabularies may be purely conventional or arbitrary. Then rote learning is the only means for building up the translation dictionary, and if the correct translations must also be discovered, immense amounts of trial-and-error search may be required.

The aspects of the environment with which we, as organisms, deal effectively reach neither of these two extremes. The translation between the state language that describes our perceptions of the world and the process language that describes our actions on the world is reducible to no simple rule, but it is not, on the other hand, arbitrary. Most of our skill in dealing with the environment is embodied in elaborate heuristics, or rules of thumb, that allow us to factor, approximately, the complex perceived world into highly simple components and to find, approximately and reasonably reliably, the correspondence that allows us to act on that world predictably. This is the skill that the adult businessman uses when he makes a decision, the skill of the scientist in his laboratory, the skill of the subject in a problem-solving experiment, the skill of a child learning to speak.

What we have proposed this evening is that at the core of these heuristics--the portion that is not bound up in special skills--is the organized system of ends-means processes, of state-process translations, that the General Problem Solver describes. We have proposed that here, in Mr. Pitts' words, is a first approximation to "the hierarchy of final causes traditionally called the mind."