

not merely a single word but a group of words.

The fact that the models must handle arbitrary environments means that the generalization from single instances to general class statements are bound to be wrong some of the time, just as inductions in general can never be guaranteed right. Thus there must be some method for unlearning, such as is found in the third program.

These models develop their "graphs of experience" (to use C. S. Peirce's phrase) in much the same manner as the human learner. Only pieces of information, often erroneous or uninterpretable, are presented. If a "teacher" (in the form of either a culture, a parent, a peer, or a formal school system) reorders the pieces into a sensible and sensitive training sequence, learning is expedited. The "wolf child" learns very little.

These models were programmed for digital computers because these computers are accessible, and extremely convenient as tools. They are actually specifications for analog devices. The analog that would embody them most efficiently would appear to be quite plausible as a brain model. Essentially, it would consist of a construction apparatus that continually builds and tears down the graph of experience, by raising and lowering the thresholds of connections between different nodes. It could, for example, be a network of neurons.

When we consider the logic of the situation, the models seem relatively plausible. Living brains have evolved as a function of rewards for responding appropriately to environmental stimuli. We are asking these models to do the same. It seems reasonable to assume that the demands of the evolutionary process are a strong constraint toward the development of the simplest (in the sense of fastest and smallest) brain that can handle the new problem confronting the organism. This is a step-by-step simplicity, where a brain that has evolved to handle one set of problems must now improve sufficiently to handle some new problem, and is thus a function of the sequencing of specific problems to the organism, and not of the entire set of problems taken at one time.

The position of the models described in

this paper is much the same vis-à-vis their environments. If models are developed that are efficient in that they can keep the graph of experience simple in the above sense, there seems reason to hope that they will be similar to living brains.

#### REFERENCES

- Richards, I. A. *German through pictures*. New York: Washington Square Press, 1959.
- Stevens, Mary E. A survey of automatic reading techniques. *Report No. 5648*. Washington, D. C.: National Bureau of Standards, 1961.
- Uhr, L. The development of perception and language. In S. Tomkins and S. Messick (Eds.), *Computer simulation of personality*. New York: Wiley, 1963. Pp. 231-266. (a)
- Uhr, L. Pattern recognition computers as models for form perception. *Psychol. Bull.*, 1963, 60, 1, 40-73. (b)
- Vossler, C., & Uhr, L. Computer simulations of a perceptual learning model for sensory pattern recognition, concept formation, and symbol transformation. *Proceedings of the IFIP Congress*. Amsterdam: North Holland Publishing Co., 1963. Pp. 413-418.

(Manuscript received November 1, 1963)

#### Argus: An Information-Processing Model of Thinking

Walter R. Reitman, Richard B. Grove,  
and Richard G. Shoup  
*Carnegie Institute of Technology*.<sup>1</sup>

Argus is an information-processing system consisting of a central executive and a network of semantic elements. Many of these elements may be active simultaneously. The executive is related to the sequential processing organization used in the General Problem Solver (GPS) (Newell, Shaw, & Simon, 1960). The semantic elements are derived from and functionally similar to the cell-assemblies conjectured by Hebb (1949).

The program described below is designed to explore problem-solving in such a system. It is also a first step in the evolution of a more general conceptual tool. We anticipate that later versions of the program will be

<sup>1</sup> The authors would like to acknowledge their indebtedness to Allen Newell, Thomas Standish, and Steven Vandenberg for their very helpful comments, and to the Carnegie Corporation of New York for its encouragement and financial support.

able to give an integrated account of cognitive structure and of such processes as understanding, cognitive learning, and a range of types of thinking extending from controlled, goal-oriented problem-solving to daydreaming and free association (Reitman, 1964a, in press).

Several other points may be noted, though they are only indirectly connected with the use of the program as a model and conceptual tool. For example, though Argus and GPS hardly span the range of possible information-processing systems, taken together they do suggest the variety of fundamentally different systems of psychological assumptions that may be represented in list-processing languages and in computer models generally. Then, too, since Argus is partly modeled after GPS and partly a reaction to the behavioral consequences of certain of its organizational concepts, what follows should indicate some of the ways in which one information-processing model may influence and affect the development of others. Finally, because Argus is intended less to simulate behavior than to explore the implications of a system of assumptions about cognitive organization and activity, it is relevant to current discussions concerning the range of uses to which such models may be put (Reitman, 1964b, in press).

We begin with an informal treatment of some of the psychological problems motivating work on Argus. After tracing out some of its connections to Hebb's theory, we then describe its major structural and processing characteristics. The paper concludes with a brief discussion of some current lines of work aimed at enabling Argus to account for more of the very impressive organizational and control abilities we find reflected in protocols of human cognitive activity.

#### **Psychological considerations: distractability and context change**

Probably the most significant work of the past decade on the processes underlying goal-directed thought is that of Newell, Shaw, and Simon (e.g., 1958, 1960; Newell & Simon, 1961), Gelernter (1960), and their associates. These investigators developed and applied information-processing models

to research in this area, and they are mainly responsible for the recognition which these new conceptual and methodological tools have achieved (Feigenbaum & Feldman, 1963; E. Hunt, 1962; J. Hunt, 1961; Miller, Galanter, & Pribram, 1960; Tonikins & Messick, 1963).

Given the need to justify and explain at the same time some particular theories and a complex and unfamiliar methodology, these earlier papers understandably focus for the most part upon the substantial similarities to be demonstrated by comparisons between human performance and the behaviors generated by the models. As a number of excellent recent discussions (e.g., Neisser, 1963; Newell, 1962b) make evident, however, the *differences* discovered through such comparisons may be equally significant keys to the underlying organization and structure of psychological activity. The program reported below arose in part in response to just such a pattern of informally observed differences—differences, incidentally, similar at many points to those noted and discussed by Neisser and Newell.

It is difficult to study the behavior of GPS without being struck by its singlemindedness. So total a lack of distractability might perhaps occur in some perfect exemplar of the Pavlovian "strong nervous system," but it is by the same token quite atypical of human thinking generally. Rarely are we capable of such utter tenacity.

Consider a few examples. Suppose the phone rings while you are writing a letter or working out a lecture. You answer, discuss some unrelated subject, and then return to your task. Only infrequently will interruptions and interpolated activity of this sort have any serious effect upon your work.

Many of us with "weaker" nervous systems can be distracted by much less insistent stimuli—for example, by the occasional unexpected ideas that crop up in the course of our work. But even if a new thought arising in this way is carried forward for a bit, we generally encounter little difficulty in recovering the strands of our previous activity, particularly if we have taken down a note or two before following out the interrupting idea. Now and then, of course, we are *not*

able to recover our previous context and at such times we may wonder unhappily whether we have not perhaps lost some great thoughts in the shuffle. But such outcomes are infrequent. As casual observation of any luncheon conversation will testify, the limited extent of such losses and confusions is amazing when we consider the range and frequency of the distractions orderly cognitive activity is subjected to.

Nor will anyone who has ever been really worried about some problem underestimate the difficulty involved in approximating the singlemindedness we see in GPS. It is no matter of closing one's door, unplugging the telephone, and concentrating on the day's agenda. Both normal and neurotic experiences attest to human inability to ignore anxiety stimuli or control their propagation. Over and over again, they return to interrupt. Quite apart from matters connected with a primary somatic signal such as pain, the recurrence of insoluble problems and their unwanted intrusions into remote cognitive contexts attest to fundamental differences between human information-processing and that which goes on in strictly sequential, centralized systems such as GPS.

### Loss of information

Finally, we would suggest that human ability to endure interruptions with a minimum of confusion is related to a more general insensitivity to loss of detail. Current computer programs typically presume perfect recall of a great deal of specific information. Errors or omissions at quite trivial levels can result in total derailment. With each subgoal it generates, for example, GPS stores context information requiring some two hundred words of code. This is the information it needs if it is to operate meaningfully with that subgoal at any later time (Newell, 1962a, 1962b). Note that these subgoals have as their terms the abstract, sensorily empty expressions of symbolic logic. Empirical problems involving sense data no doubt would require a great deal more stored information. Like Tristram Shandy, such a system might soon spend much of its life recording the details of a single day. Humans, by contrast, seem gen-

erally less dependent upon a multitude of details, more able to regenerate what they need as the occasion arises. One might well argue that as a general rule, almost everything we experience tends to be forgotten. Only as an exception is an item of information long retained and ready at hand.

To summarize, we are subject to a variety of internal and external distractions, and we readily lose information. Nonetheless, we continue to think and to solve problems, and with much less difficulty than observation of current information-processing models might lead us to expect. The discrepancy raises several questions. What, for example, are the arrangements mediating interruption? What determines our choice of one problem over another if we are distracted and confronted with alternatives? How can we explain our relative insensitivity to loss of detail?

We might account for our ability to handle telephone interruptions by analogy with mechanisms now utilized in many computer systems. Interruptions by worry or promising ideas are another matter, however. To explain them, we must assume a system sensitive to its understanding of, and interest in, the information it is processing, and capable of assessing the likely importance of incipient alternatives. If we are to account for its imperfect ability to put aside problems and ideas as it chooses, the system also must be assumed to include a *number* of parallel centers of activity, rather than just one as in GPS. Finally, we must specify how choices among alternatives might be determined by relative significance or urgency, or perhaps even by the relative sensitivity to disturbance of the information contexts serving to frame and define each of the alternatives. Available information-processing models enable us to pose and to think about such questions, but they are not a sufficient basis for dealing with them.

### Alternatives to strictly sequential, centralized processing

In his Introduction to *The Organization of Behavior*, Hebb labels "the failure . . . to handle thought adequately" as "the essential weakness of modern psychological

theory" (1949, p. xvi). Defining thought as "some sort of process that is not fully controlled by environmental stimulation and yet co-operates closely with that stimulation" (p. xvi), Hebb then develops a schema intended to serve as "a conceptual tool for dealing with expectancy, attention, and so on, and with a temporally organized intracerebral process" (p. xviii). The result, built around the cell-assembly concept, is very likely the most seminal systematic theory in psychology which appeared in the postwar era, and the idea of a system of active interacting elements is very suggestive when viewed against problems of the sort we have been discussing.

Substantial as its other merits may be, however, in no sense can Hebb's theory be considered a general theory of thought. In particular, Hebb never shows how his theory might account for goal-directed thinking, e.g., as one finds it reflected in protocols of problem-solving behavior. Nor have we been able to imagine any way in which a system consisting entirely of Hebbian cell-assemblies might be made to do so. "Thought," in fact, has only one index entry for the entire body of Hebb's book. "Thinking" and "problem-solving" get none at all. Thus it is not surprising that though there have been computer models of information-processing based upon Hebb's theory, they have for the most part been concerned with the development of cell-assemblies from individual neurons (e.g., Rochester, Holland, Haibt, & Duda, 1956) rather than with the processes of thought.

There is another difficulty involved in simulating thinking in a system of active parallel elements of the sort Hebb describes—the very vagueness of the concept of a parallel system itself. Just as the assertion that a system is nonlinear tells us nothing about the kinds of deviations from linearity it shows, so terming a system parallel implies nothing about the extent of parallel activity or the constraints to which it is subject. By presenting evidence bearing on things humans can *not* do, Broadbent (1958, 1962), Miller (1956), and others have suggested certain limitations that parallel models of human information-processing would appar-

ently have to reflect. Within these limitations, however, many organizational principles remain which might serve as bases for explanation of those performances humans *are* capable of: with respect to selections from among these, psychological theory in general seems to have very little to say.

In his paper on the organization of GPS, Newell (1962b) points out that "sequential processing . . . encourages us to envision isolated processes devoted to specific functions, each passively waiting in line to operate when its turn comes. It permits us to think of the total program in terms of only one thing going on at a time" (p. 398). Without a centralized, sequential arrangement, he suggests, information-processing might well be like an "Alice-in-Wonderland croquet game, with porcupine balls unrolling themselves and wandering off, and flamingo mallets asking questions at inopportune times" (p. 398).

Nonetheless, though the over-all organization of GPS with its goal structures and associated context storage procedures clearly is an excellent basis for carrying on complex information-processing in an orderly fashion, discussions such as Hebb's and the informal examples of human experience considered above convince us that human cognitive activity must be organized at least in part in a fundamentally different framework. Though we cannot see how cell-assemblies and phase sequences *by themselves* can be made to provide an adequate general account of thought, perhaps a system linking a limited sequential control to an underlying structure of active elements fashioned after the Hebbian model might furnish the necessary basis.

Argus thus is a first step toward a system that begins by positing active cognitive elements and gradual loss of information. It is being used to explore how an organism operating under these conditions might be enabled to get back to the right track close enough and often enough to achieve some semblance of progress in its work. In Newell's metaphor, in other words, Argus is an inquiry into the virtues of flamingoes given a modicum of order in Wonderland.

### Program structure and operation

The program described below is written in IPL-V (Newell, 1961) and has been run on the Bendix G-20 at Carnegie Institute of Technology. As was noted above, Argus is written so as to make possible extended versions applicable to a broad range of cognitive functions. The versions now running, however, are limited to analogies problems of the form  $A:B :: C:(W, X, Y, \text{ or } Z)$ . To simplify the exposition, minor differences among the several versions of Argus are for the most part ignored here. All versions share the same basic organizational structure, and each is capable of finding solutions to the simple analogies problems used in testing the system.

After the main components are described, some of the dynamics of the program in action will be discussed. Since program development is continuing, specific details are given only when they are necessary to an understanding of the basic organization of the system.

The main aspects of the system are the sequential *executive*, the *network* of active semantic elements, and the channels of *interaction* between them. There are also *experimenter* routines, but since these are concerned primarily with input and output, performance monitoring, and similar house-keeping tasks, they need not be described further. The executive and the routines for simulating the parallel activity of the semantic elements together require somewhat more than one thousand IPL-V words. The number of words required for the network varies, of course, with the number of its elements and the extent of their interconnections.

### Executive

The integrative activity of the executive is organized on four levels. From the top down, these are: subject, problem, strategy, and strategy step. We may think of an individual as someone who has available a variety of information-processing strategies along with some rules for their use in specific situations. An individual subject is therefore defined as a set of information-processing strategies and rules for their use. Given a

subject, that is, a particular set of strategies and rules, the Argus experimenter generates a problem. All problems are analogy items at present, and each is read in as a list of seven words corresponding to the names of the three given semantic elements and the four answer alternatives. Using the strategies and associated rules specified, the executive now tries to solve the problem. Each strategy is a scheme built up of steps, the basic units of behavior available to the executive. Figure 1 depicts a simplified version of one such strategy. Though rudimentary and relatively ineffective as given here, the strategy illustrates the kinds of unitary behaviors assumed and the ways in which they may be organized in higher-order information-processing schemes. The implications of the various steps will be evident once we have discussed the structure of the semantic elements upon which they operate, and are ready to consider the executive cycle and the dynamics of the Argus system as a whole.

Changes in the semantic network occur *within* steps. When a step commences the network is in some state  $s$ . By the end of the step, changes may have occurred in the status of individual semantic elements, and in the strengths of interrelations among the elements. Thus the network as a whole is in some new state  $s'$ . No further changes take place until the next step begins. The order in which changes within a step occur is only locally relevant. Such changes become available to the executive *together*, and only after the step in which they occur has been completed. The step thus serves a double function. Like the method segment in GPS, it is the minimal unit of behavior. In an important sense, however, it is also the minimal unit of time; from the viewpoint of the executive, network changes occurring within the step may be thought of as taking place in parallel and at approximately the same time.

### Cognitive structure as a network of active semantic elements

A detailed discussion of the representation of complex cognitive structures by semantic networks is available elsewhere (Reitman,

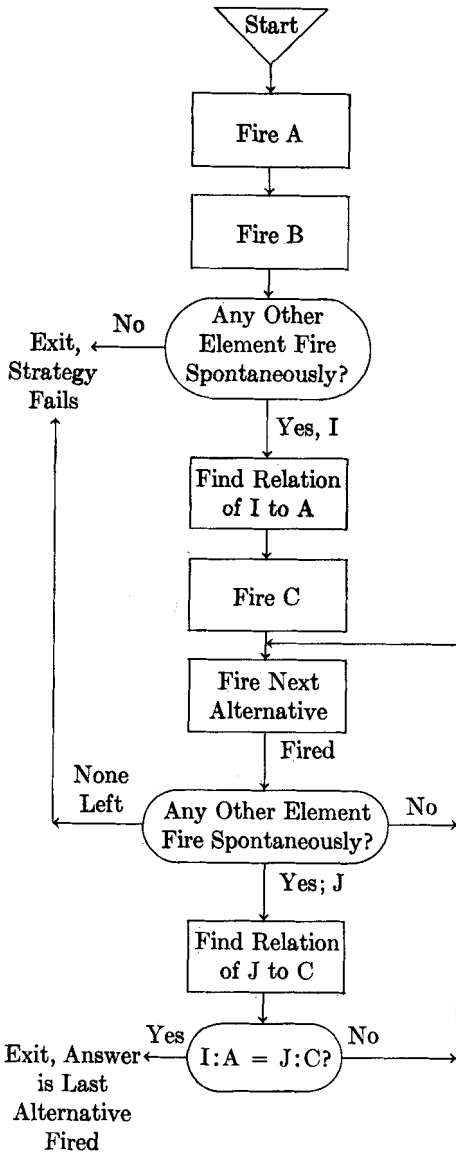


FIG. 1. Rudimentary Strategy (Slightly Simplified) for Analogy Problems of the Form  $A:B :: C:(W, X, Y, \text{ or } Z)$ . When an element is fired, activation and inhibition increments are added to related elements, altering their states.

in press). Here we consider only those aspects of the Argus program that enable us to realize such a network under the additional assumptions that the network elements are active, and that many interactions may go on among them simultaneously.

The basic units of the Argus cognitive structure are the semantic elements. These

correspond to Hebb's cell-assemblies in terms of the psychological functions they serve. Each semantic element is an IPL-V structure consisting of three description lists. The first two lists define the direct connections or associations to other semantic elements in the cognitive structure. Since Argus at the present time has neither perceptual inputs nor behavioral outputs, the meaning of a semantic element is solely a function of the set of its direct and indirect interconnections to the other elements of the cognitive structure.

Each relation to another semantic element is specified by an attribute-value pair on the first description list. For example, the relation of the semantic unit corresponding to the sensation of heat to the one representing its opposite, cold, would be given by the pair (OPPOSITE, COLD) on the first description list of the semantic element HEAT. The pairs on the second of the three lists defining an element specify the current strength of association for each such relation between the element and another semantic element in the cognitive structure.

The third description list gives the parameters defining the current state of the semantic unit. These parameters are named in terms reminiscent of Hebb's conceptual nervous system. This is an informal usage, however, primarily for mnemonic purposes, and carries no neurophysiological implications about corresponding elements in the human nervous system. That is, the Argus semantic elements are purely psychological or information-processing constructs. They correspond to units of meaning in human cognitive structures, and their existence is taken as given.

Five parameters are required to specify the current state. The first four give activation, inhibition, threshold, and size values for the semantic element. The fifth names the internal date when these values became current. Activation and inhibition values decay with time, and all of the first four parameters may vary with other changes occurring in the system. Thus the values recorded for these parameters, though correct as of the indicated internal date, will in general be wrong when the semantic ele-

ment is called or acted upon subsequently. Any routine affecting or using a semantic element must therefore first call in sub-routines that compute appropriate current values for the first four parameters and update the fifth.

**Activation, inhibition, and conflict.** It might seem simpler to replace activation and inhibition by a single parameter representing the net resultant. The modification certainly would be advantageous computationally. We maintain the distinction because we are working at an aggregate level, i.e., at the level of semantic elements, rather than with units corresponding to individual neurons. It is not obvious that we want to consider a semantic element which is high both in activation and inhibition to be psychologically equivalent to one having low values on these two parameters. On the contrary, semantic elements high both in activation and inhibition provide a very reasonable representation of one type of cognitive conflict, i.e., a conflict focused about a few specific elements of meaning in the system. As we shall see in what follows, the occurrence of such a conflict anywhere in the cognitive structure is one of the several kinds of spontaneous events taking place there that may be noticed by the executive and that thus may have substantial effects upon the subsequent behavior of the system.

**Threshold, size, and propagation of activity.** The first two description lists of a semantic element specify the elements it is associated with and the strength of each of these associations. The set of all such pairs thus describes the structure of interconnections within the cognitive structure as a whole. Following Hebb (1949, p. xix), we think of the semantic elements in Argus as generally open in the sense that usually changes in the state of an element might be expected to have some effects beyond its bounds. But there seems to be no good basis on which to specify the extent or significance of these constant low-level interelement interactions. Worse yet, if they are extensive and frequent enough, there is no way in which a serial device can be expected to continue to simulate more than a small frac-

tion of such interactions—the Tristram Shandy problem again. A second problem arises out of Hebb's conception of the firing of an element once activity within it rises above a certain level.

In Argus, both problems are handled by a single parameter, the threshold. Thus this concept in Argus is as much a programming convenience as it is a theoretical construct. Threshold marks the point of maximum activity within an element. Below threshold, activity within an element does not propagate. Once activity reaches threshold level, the element is said to fire.

The maximum impact any particular element may have on others is specified by its size. Size is thus an ancillary parameter, reflecting the assumption that some elements are "larger" or otherwise more important than others in determining the direction of activity within the cognitive structure.

We have stated that propagation takes place within a step. The current step terminates only when no further elements are active enough to fire. Interminable firing loops are prevented by conventions, e.g., by permitting elements to fire only once within a step, or by reducing the activation level of a just-fired element to some arbitrary proportion of its threshold value. The latter procedure to some extent resembles the refractory period behavior of individual neurons, of course, but again, since the semantic element is an aggregate construct, the force of the analogy is uncertain and the procedure may better be regarded as a pure convention.

#### **Interactions between the executive and the active semantic network**

**The signal system.** One of the main channels for interaction between the executive and the semantic network is a system of signals and signal cells based on that used in GPS. During a step, signals may be set for any of several reasons.

As we indicated above, the occurrence of a cognitive conflict anywhere in the semantic network is one of the several kinds of spontaneous events taking place there which may be noticed by the executive and which thus may have an effect on its subsequent

behavior. Specifically, whenever the activation and inhibition levels for an element both exceed an arbitrary limit, a conflict signal is entered in the main signal cell and the name of the element is recorded elsewhere in the signal system as the immediate locus of the conflict. Now the executive is in a position to notice this conflict upon completion of the current step. What the executive does at that point will depend among other things upon the characteristics of the current subject and current strategy. Some people are easily distracted from a line of thought by such conflicts. Others may attend to them selectively, depending upon the locus of the conflict and what they are doing at the time. Still others may ignore them, or perhaps not even be aware of them. The use of separate activation and inhibition parameters permits us to represent such conflicts about a semantic focus, and it will make possible detailed exploration of the problem-solving consequences of individual differences in strategies for reacting to them.

Similarly, whenever an element fires as a result of increments in activation propagated from related elements, a signal is entered in the main signal cell and the name of the spontaneously firing element is recorded in an ancillary cell. Once again, any of a variety of responses may be made to such an event once the current step is completed.

The main signal cell is a pushdown list and thus may accumulate any number of such signals in the course of a step. In addition to those initiated by cognitive conflict or spontaneous activity, it may contain signals set with or by the strategies, and by other sources as well.

**The executive cycle.** Once a step has been completed, responsibility for decisions about further action passes back up to the executive. The executive checks the main signal cell and then makes its decision in the light of the signals it finds there and the dictates of its current strategy. Thus the signal system makes possible a quite general branch upon the state of the organism as a whole, including some aspects of activity in the cognitive structure for which the executive is not in any direct sense responsible.

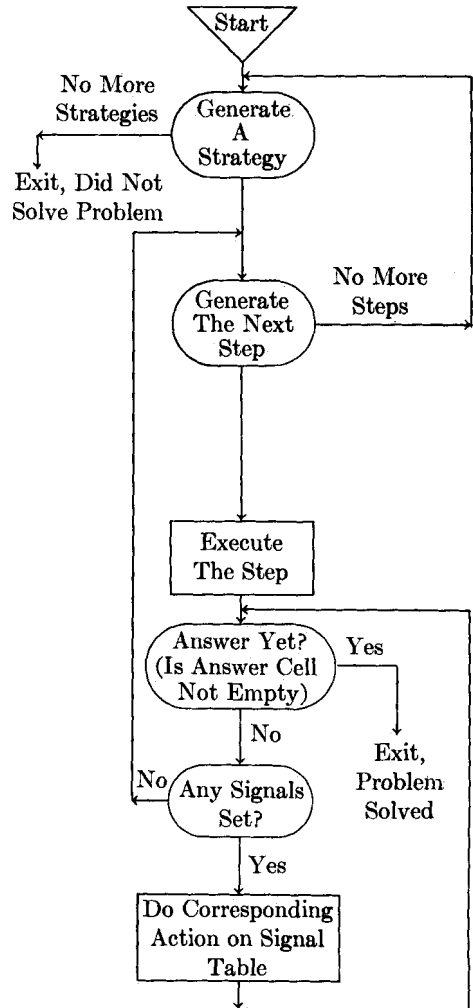


FIG. 2. Executive Cycle (Slightly Simplified). Subject and problem are specified beforehand. This gives the executive a set of strategies and rules for their use. The signal-action pairs in the signal table at any given time are a function of the rules.

This generalized branch is achieved by means of the executive cycle, depicted in slightly simplified form in Figure 2. Though not at all complex in itself, the executive cycle makes possible a very high degree of flexibility for the Argus system as a whole. Specifically, as Figure 2 indicates, once a step is executed, if no answer has been obtained to the current problem, the executive checks the signal cell. If there are no signals, it simply continues with the next step of the current strategy. If one or more signals are

present, however, the executive's next action is taken from a current signal table, and its behavior is thus a joint function of the signals present and the particular subject and strategy under which the executive is operating. As alternatives to continuing to the next step in the current strategy, the executive may choose to jump to a nonsequential step in the strategy, or to another strategy altogether. The housekeeping routines associated with the system are such that there are no limitations on the executive's choice at this point.

Note that this indirect mode of operation through the signal system means that actions need not be invariant with particular signals. Furthermore, there is a general reset routine which modifies the table of signal-action pairs as a function of information taken from the description lists associated with the current subject and the current strategy. In fact, one of the actions open to the executive in response to certain signals is just such a resetting of elements of the signal table.

It is very difficult to communicate the possibilities that an arrangement as flexible as this opens up. Certainly it should be clear, however, that the execution of a strategy, even one as rudimentary as the one depicted in Figure 1, is in no sense a cut-and-dried affair. The modulation of behavior in response to signals permits events in the semantic network to exert a pervasive influence upon the over-all sequential control of the system. As a result, Argus now manifests something approaching the distractability and interruptability discussed in detail at the beginning of this paper.

**Executive controls over activity in the semantic network.** We have now described the structure of the executive, the structure and dynamics of the semantic network, and some of the conditions under which events in that network affect problem-solving behavior. There are two principal means whereby the executive may in turn affect events taking place in the semantic network. The first of these is by the set of search, interrogation, and generation procedures programmed into the strategies (see Figure 1 for examples). In general, when Argus elements are stimulated as a consequence of

such executive action, their activity levels and effects upon related elements increase. Some of the strategies in fact deliberately instigate such activity in order to observe the subsequent firing.

A second means of control is possible in current Argus programs, though none of the strategies written so far make use of it. This would involve manipulation of the activation or inhibition levels generally or in portions of the semantic network. There is excellent experimental evidence that such level-changes take place in humans, and the experiences of relaxation and concentration might be interpreted in terms of them. In addition to its possible psychological significance, such local manipulation of levels is of interest as a potential control device for information-processing in artificial intelligence systems. In either case, the result is that Argus has at its command a probe of variable width and intensity which is capable of everything from gross discriminations to the finest and most precise effects, very much like that variable beam of light with which consciousness and attention have sometimes been compared.

### Current developments

Several reasons were presented at the beginning of this paper for exploring alternatives to strictly sequential systems such as GPS. These included human sensitivity to distractions from without and from within, and human ability to re-establish appropriate contexts after such interruptions, often despite intervening loss of detail. In plans for development of the Argus programs, two general methods have stood out which might enable a system of this sort to restore or rediscover context without explicitly preparing for context changes as GPS does. The first of these is a time line; the second is a system for generating and using internal sentences. Each exists in preliminary form in current running programs. Since a great deal of work remains to be done on both, however, the present discussion is intended largely to suggest something of the roles these developments are expected to play.

**The time line.** The basic idea for the time line comes from Penfield and Roberts (1959), whose experimental stimulation of human cortex *in vivo* produced results which they interpret as suggesting that humans lay down a temporally-ordered cortical record of their experiences. Several Argus programs also lay down such a record. In later versions of Argus, this sequence of semantic elements corresponding to a sequence of experiential events will serve several important purposes. First, though incomplete cues may by themselves be inadequate as a context for problem-solving, when stimulated they may propagate activity to (and fire) the time-line record to which they refer. With such an entry to a relevant point in the time line, the problem-solver may have access to enough information on the original context at that point in the time line to be able to restore or reconstruct the balance.

The use of a time line also makes possible a rather straightforward and quite general treatment of forgetting. Since the time-line entries are themselves semantic elements, and since all associations among elements decay with time, no specific erasure of information is necessary to account for the loss of transitional objects created during information-processing, or for the forgetting of previous events in general. The only time-line entries which will continue to be retrievable will be those strongly tied in with other elements in the cognitive structure, e.g., as a result of frequent reference. As their fore and aft connections within the time line decay to a low level, these individual entries will in effect come to be ordinary semantic elements, representations of the significant experience of the system. As such, they continue like other elements to be available to the executive, perhaps to be retrieved, for example, in the course of a search for specified concrete instances to use in deriving or testing some generalization or abstraction. In this way, moreover, the time line becomes an important source of new lexicographic units for the system, units which may enter into new attributes employed to describe still other objects (Reitman, 1964a). Thus, at least one kind of learning will be made a continuous and integral response of Argus,

in keeping with the guiding aspirations described at the beginning of this paper.

The studies of Penfield and his associates, a number of psychoanalytic studies, and certain investigations utilizing hypnotic techniques all quite strongly suggest that the details of many experiences which no longer can be retrieved by strategies normally available to the executive nonetheless remain intact in memory and may be evoked by suitable techniques. In Argus, such a phenomenon could be reproduced by means of the time-line device with a very simple adjustment of the decay curve for associations. Alternately, since this phenomenon is not expected to be of practical significance for the sort of problem-solving behavior in which Argus will engage, the system may include an infrequent sweep through memory which returns linking cells to available space when the association between elements falls below some minimal strength.

Exploration of the time-line device suggests interesting alternative ways of looking at a number of traditional information-processing questions in psychology. To cite just one example, the lines between recalling, reconstructing, and recreating information become much less clear. Suppose that the organism executive has gained access to one event in the time line and tries by direct stimulation to activate records of immediately subsequent events (i.e., "tries to remember what happened after that"). Should that fail, the executive might switch strategies and try instead to reconstruct what might have happened next, from its knowledge of the goals and status of the situation at the event in the time line accessible to it. Suppose, furthermore, that now an immediately subsequent entry in the time line fires, due to activation propagated to it from elements stimulated during the executive's attempts at reconstruction. It examines the newly retrieved element, tests for its relation to the previously available event record, just as it would have done had an event record fired under the direct recall strategy, and concludes something equivalent to "Oh yes, I remember." But in this view, the main difference is at the strategy level, at the level of intent. In terms

of the pattern of activation, it might well be impossible to separate those contributions to activity pattern changes due to "recollection" from those due to "reconstruction" or "imagination."

**Internal sentences.** One need only overhear a snatch from a conversation—e.g., "Mantle hit four runs. . . ." to realize how quickly humans can sometimes establish an appropriate context. Although the means whereby humans achieve the result remain unclear, such programs as Lindsay's (1963) provide some quite useful concrete suggestions. Current efforts toward an internal sentence mechanism for Argus are predicated on the assumption that the power to reproduce an appropriate context from a sketchy sequence of clues accessible in memory may very well depend on the same mechanisms that enable humans to understand equally sketchy external sentences, presumably by calling forth a context into which the information they convey is fitted. In conjunction with the time-line device, such a mechanism might provide an inexpensive and yet extremely powerful means of accounting for the insensitivity to loss of detail and the ability to restore or reconstruct context attributed to humans at the beginning of this paper.

Anyone who has had occasion to compare the comprehensibility of his lecture notes after a day, a week, and a year is familiar with the way in which such material grows cold. Archivists experience a somewhat similar phenomenon on a longer time scale. The design of library retrieval systems is complicated by a tendency for key terms to change in meaning over the years. Older classification schemes lose relevance to current usage, and it becomes increasingly difficult to retrieve pertinent older material unless one knows the meanings the category labels carried at the time the material was classified and stored. It is to be expected that a system built around time-line entries and internal sentences will show these same effects as some lexicographic units decay out of reach or are replaced by others, and as the relation-association-strength pairs change values with time. Such a result might perhaps be a drawback for some artificial intelligence applications, but it seems a clear

advantage for a model of human information-processing which aims to derive such phenomena as learning and forgetting from a small number of simple and universally applicable system axioms.

**Generalizing the cognitive structure.** Perhaps we may conclude by noting one other longer-range aspiration. As things stand now, a certain amount of the "intelligence" of the system is encoded in special formats, notably the signal table and the strategies themselves. We would like to experiment with having the functions these components serve taken over by the general cognitive structure, so that all that remains of the executive is an organizational focus, something insuring the "modicum of order" we were unable to account for with the Hebbian model. Like the Marxian state, the sequential executive may prove resistant to attempts to let it wither away, but the possibility has a theoretical simplicity and parsimony about it that make it very attractive.

#### REFERENCES

- Broadbent, D. E. *Perception and communication*. New York: Pergamon, 1958.
- Broadbent, D. E. Attention and the perception of speech. *Sci. Amer.*, 1962, 206, 143-151.
- Feigenbaum, E. A., & Feldman, J. (Eds.). *Computers and thought*. New York: McGraw-Hill, 1963.
- Gelernter, H. Realization of a geometry-theorem-proving machine. *Proceedings of the International Conference on Information Processing*. Paris: UNESCO, 1960. Pp. 273-282.
- Hebb, D. O. *Organization of behavior*. New York: Wiley, 1949.
- Hunt, E. B. *Concept learning: An information processing problem*. New York: Wiley, 1962.
- Hunt, J. McV. *Intelligence and experience*. New York: Ronald Press, 1961.
- Lindsay, R. K. Inferential memory as the basis of machines which understand natural language. In E. A. Feigenbaum and J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill, 1963. Pp. 217-233.
- Miller, G. A. The magical number seven, plus or minus two. *Psychol. Rev.*, 1956, 63, 81-97.
- Miller, G. A., Galanter, E., & Pribram, K. H. *Plans and the structure of behavior*. New York: Holt, 1960.
- Neisser, U. The imitation of man by machine. *Science*, 1963, 139, 193-197.
- Newell, A. (Ed.). *Information Processing Language-V manual*. Englewood Cliffs, N.J.: Prentice-Hall, 1961.
- Newell, A. A guide to GPS-2-2 program. Pitts-

- burgh: Carnegie Institute of Technology, 1962 (ditto). (a)
- Newell, A. Some problems of basic organization in problem-solving programs. In M. C. Yovits, G. T. Jacobi, and G. D. Goldstein (Eds.), *Self-organizing systems*, 1962. Washington: Spartan Books, 1962. Pp. 393-423. (b)
- Newell, A., Shaw, J. C., & Simon, H. A. Elements of a theory of human problem solving. *Psychol. Rev.*, 1958, 65, 151-166.
- Newell, A., Shaw, J. C., & Simon, H. A. Report on a general problem-solving program. *Proceedings of the International Conference on Information Processing*. Paris: UNESCO, 1960. Pp. 256-264.
- Newell, A., & Simon, H. A. Computer simulation of human thinking. *Science*, 1961, 134, 2011-2017.
- Penfield, W., & Roberts, L. *Speech and brain mechanisms*. Princeton: Princeton Univ. Press, 1959.
- Reitman, W. R. Heuristic decision procedures, open constraints, and the structure of ill-defined problems. In M. W. Shelly and G. L. Bryan (Eds.), *Optimality and human judgment*. New York: Wiley, 1964. (a)
- Reitman, W. R. Information processing models in psychology. *Science*, 1964, 144, 1192-1198. (b)
- Reitman, W. R. *Cognition and computer simulation* [working title]. New York: Wiley, in press.
- Rochester, N., Holland, J. H., Haibt, L. H., & Duda, W. L. Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Trans. on Information Theory*, 1956, PGIT-2, No. 3, 80-93.
- Tomkins, S., & Messick, S. *Computer simulation of personality*. New York: Wiley, 1963.

(Manuscript received March 27, 1963)

## Natural Language Inputs for a Problem-Solving System

Gilbert K. Krulee, David J. Kuck, Dale M. Landi, and Denis M. Manelski  
The Technological Institute, Northwestern University, Evanston, Illinois.<sup>1</sup>

### INTRODUCTION

When we ask students to solve problems or prove theorems, we usually communicate to them in English and take for granted their

<sup>1</sup>This research is supported through Grant G-17951 from the National Science Foundation, and is part of a larger program of studies on the simulation of human behavior. The design of the system has been substantially completed. The system is currently being programmed using the IPL-V list processing language (Newell, 1961).

ability to comprehend what has been written or said. In communicating with a digital computer, we expect difficulties and customarily begin with a highly formalized problem statement which can be processed directly by the computing system.

The research described in this report explores methods for more direct and natural communications between operator and computer. We have chosen for study a limited class of problems: those in the propositional or statement calculus. Our objective is a computer system that will accept problem statements in English and be able to "understand" these natural language inputs.

In the design of our system, we have been very much influenced by "Baseball," a computer program written by Green, Wolf, Chomsky, and Laughery (1961). The structure of their program may be summarized as follows. "Baseball" answers questions about baseball which are stated in ordinary English; it is essentially an information retrieval program with two major subsystems. The first is primarily linguistic. It analyzes a question, using some knowledge of English syntax, and produces as its output a specification list summarizing the meaning of the question. This list is then used as input to the processing system, which proceeds to answer the question as best it can. As the processing system searches among its store of available data, its search procedure is governed by the details of the specification list. In a sense, this list provides a formal representation for the problem. From a logical point of view, the formal representation is not unlike an equation with an unknown (the question) and a set of conditions (the information given) which must be satisfied in order to reach a solution (Stoll, 1961, p. 22).

Our logic system is also separable into two major parts, the first of which is linguistic and accepts as its inputs a set of statements in English. The output of this "extraction" subsystem is a formal specification, which, in turn, is used as input to the problem-solving portion of the system. These are in the form of a set of premises and a conclusion. It is the task of the problem-solving system to verify this conclusion. The feasibility of the problem-solving portion of this system has been demonstrated elsewhere (Newell &