SYSTEM RESEARCH LTD

R. T. Mckinnon Wood MA Gordon Pask MA PhD S. M. Mckinnon Wood Elizabeth Pask BA

A non-profit organisation for research in cybernetics and the behavioural sciences

20 HILL RISE · RICHMOND · SURREY · Telephone : RIChmond 0801

AFOSR 67-2001

CYBERNETICS AND EDUCATION

Gordon Pask

Paper prepared for presentation at the
Fifth Congress of the International Assn. of Cybernetics,
Namur, Belgium, September 1967.

Research sponsored by the Air Force Office of Scientific Research through the European Office of Aerospace Research under Contract F61052-67-C-0010.

17

1. Introduction.

This paper reviews the main contributions of Cybernetics to the art and science of teaching (with emphasis upon the field of "teaching machines" and the related field of "computer-assisted instruction"). These contributions are of two sorts:

(1) Cybernetics, applied to psychology, has provided a number of useful insights into how people learn, how they are motivated, and how they direct their attention; as a result, we are better able to describe and model the learning process.

(2) Cybernetic principles have been directly applied to the prescription of teaching procedures and to an analysis of the stability of tutorial systems (in particular, man-machine tutorial systems).

In the span of this paper it is impossible to list these contributions exhaustively; instead, I shall pick out a few salient ideas and will try to illustrate their practical consequences. Amongst the important areas that are omitted from this survey are (I) The analysis of maturation, early learning and imprinting in terms of Cybernetic models (an activity of growing importance in child psychology), and (II) Cybernetic models for the social systems involved in education.

2. Learning

In the following section we set out several ways in which current theories of learning have a Cybernetic calibre or embody Cybernetic principles. Having done so, we shall examine the bearing of these principles upon teaching methods.

(1) When a man learns, he chiefly learns how to compute and control (rather than learning lists of facts or disparate

associations). So far as computation is concerned, he learns how to compute the invariant and characteristic features of a familiar face, rather than learning a list of facial details. So far as control is concerned, he learns what to do about his environment in order to achieve a goal (to carry out a controlling process of some sort) and, in the course of learning, he patches together information about how to achieve the subgoals or components of the goal that is aimed for. As a result, the conceptual descriptions which we gain by learning, serve also as prescriptions for what might be done to achieve a goal (consider, for example, how we learn to find our way home, or how we remember a play; in each case, there is an underlying plan or story which guides our learning or remembering).

(2) Prescriptions for goal achieving action may be externalised (the pilot learns to manipulate the controls of an aircraft; so as to keep on course). On the other hand, they may be internal and symbolic. The mathematician, for example, learns to mentally manipulate a mathematical structure in order to satisfy a set of abstract rules. In either case, we may regard the "control activity" as "problem solving" and consequently adumbrate our previous argument by the contention that "man learns to solve problems".

(3) More basic motivational and attention directing control systems act to determine what counts as a problem; to winnow out significant and "problem posing" events from the flux of experience. We are just beginning to understand these underlying control systems. But their existence must be respected by any competent teaching method.

(4) A large number of skills are performed in a discontinuous, partitioned, and (technically speaking) feedforward fashion.

The subject commits himself to achieving a subgoal (to recognize something in a perceptual skill, to do something in a motor skill) by triggering off a process which carries on autonomously until the goal is achieved. But even though a skill may be performed in a feedforward mode, its instruction invo s feedback from the environment*. Two types of feedback signal (delivered by the instructor) are of immediate importance; (I) "Knowledge of Results" signals that "Reinforce" the subject by confirming that a problem has been rightly solved, and (II) "Cueing Information" which simplifies the problems posed (by the instructor) and helps the student to solve them. The effect of "knowledge of results", reinforcement and cueing information, is rather well understood, at any rate for simple types of learning.

(5) Learning is an hierarchically organized process. At the lowest level we learn by constructing problem solving procedures that act upon problems posed by the environment. But this constructive activity may also be conceived as a higher level of problem solving (we learn to solve the internal problem engendered by the lack of suitable problem solving procedures) and this view of things can be respectably formulated within a theory of concept learning wherein levels in the hierarchy correspond to classes of conceptual structures. At a still higher level in hierarchy, we learn how to learn (we develop suitable attitudes towards whole classes of concepts or domains of knowledge).

* A feedforward control system contains a model that predicts the performance of the controlled element. It uses this model to determine its actions. A feedback system uses information regarding departure from the desired goal for this purpose. Many control systems use a combination of the feedforward and feedback modes.

(6) Man is not a passive responder to stimuli. He actively seeks novelty and having found novelty in the environment, he aims to make sense of it by learning how to control it. In Cybernetic jargon, "man is an active control system" who looks for new goals to achieve (this is a theoretically tractable entity). In psychological terms "man has a Curiosity Drive" (and it is one of the most important of all drives).

(7) On this basis, we may argue that man (perhaps some of the higher animals also) is impelled by a definite need to learn.

(8) Teaching is the control of learning. It is so in the following sense:

(I) From (1) and (2) above, man learns to control his universe by solving the problems that are posed to him.

(II) From (4) above, the institution of one or more new control facilities (a teaching process) entails the temporary presence of an outside control loop which delivers suitable feedback signals and catalyzes the learning process.

(III) The outside control loop is a vehicle for co-operative interaction with the student; the instructor senses the student's performance and provides him, co-operatively, with reinforcement and with sufficient cueing information to render his task intelligible at his current level of proficiency. However, the co-operation is partial only; for, to satisfy (6) and (7), co-operation must be withdrawn as the student becomes more proficient (the task must retain the student's interest and give him enough to learn about).

(IV) As stated so far, the partially co-operative interaction involves a single level of discourse. The tutorial conversation is a matter of posing problems and getting responses that designate solutions. However, (5) and (3) suggest that an acceptable teaching system should be an hierarchically organized control system rather than a single level control system; that there should be many levels of discourse in the tutorial conversation. In addition to posing problems and delivering reinforcement, the instructor in such a system makes comments about properties of the performance; in addition to responding, the student will comment upon and modify the mode of instruction.

3. Teaching

These principles can be (and have been) profitably applied to the conduct of classroom instruction and industrial training where the external control loop is completed by a real life instructor. However, it is often convenient to automate some of the functions performed by the real life instructor (that is, to replace the instructor by a teaching machine). In this paper, I shall concentrate upon cases in which this expedient has been adopted. Automation is possible just insofar as we do have a model for the learning process.

For expository purposes, it is useful to demarcate teaching that involves the inculcation of a well defined sequence of concepts (one of which must be grasped before the next is presented) and teaching where the conceptual rules are given and the student is required to gain proficiency in their application. The former type of teaching is concerned with subjects like mathematics, history, or statistics and its instrumentation calls for a sequential teaching machine. The latter is concerned with the instruction of skills (driving a motor car, using a typewriter, inspection of parts, running a chemical plant, or making inferences from a sequence of evidence). It calls for a skill teaching machine. Since the instructional framework is much simpler, it is not surprising that skill teaching machines have been built to mediate more sophisticated forms of tutorial conversation than sequential teaching machines. However, this reflects the present state of the art rather than an absolute limit (indeed, the whole distinction between one sort of system and the other, is a convenient fiction; language teaching, for example, entails the inculcation of grammatical rules as well as the exercise of these rules; ultimately, language is a skill).

4. Sequential Machines

The instructional materials are reduced to a logical sequence, in the simplest case to a "linear programme" in which one concept or problem area follows after another. To each concept, the programmer assigns one or more "frames" which are visually presented to the student. Ideally, each frame contains some expository material (explaining the concept) and it poses certain problems which are pinpointed by questions, about this or the preceding frames. The student is required to solve the problem and to make an answer. Commonly, the questions are associated with "multiple choice" answers and the student responds by selecting one of these "multiple choice" alternatives*. The programme also carries a code whereby the programmer informs the teaching machine which response is the correct response for each frame.

*Various teaching machines are available. In some of these, the student is required to write in a response rather than selecting alternatives and, if so, he is usually also required to compare his response with the "ideal response" and to take remedial action. When the onus is placed, to this extent, upon the student, the machine can be dispensed with and the programme presented as a "programmed book"

Physically speaking, the student sits in front of a screen on which the frames are displayed in sequence. For each frame he responds by pressing a response button on the machine which corresponds to one of the multiple choice alternatives. The machine compares the correct code with the student's response. If his response is correct, he is reinforced by an affirmative "knowledge of results" signal and the machine advances its programme to the next frame. The programme is written to ensure rather a high probability of correct response and it is crucial to notice that the student proceeds through the material at his own pace.

Branching programmes are more sophisticated. Like linear programmes, they determine reinforcement for each frame, but they contain more frames than are strictly necessary. A particular individual will only see some of these because any student may take several paths through the programmed material depending upon the mistakes he makes. As the name suggests, the programme contains "branch points" at which different paths may be selected depending upon the student's response.

Branch points serve a couple of different purposes. (I) They allow for remedial treatment and (II) They allow different sub-programmes of the programme to be presented to different individuals.

(I) If, at frame A in the programme, the student gives the correct response, a, he proceeds to the next frame. If he makes mistake b, he is directed through a subsequence of material (say frames B_1, B_2,............) that remedies the misconception which gave rise to b and which ultimately returns him to frame A (after which he continues with the programme). If he makes mistake c, he is led around a remedial sequence

C_1, C_2,....... and back to A. The information in the remedial sequence is essentially cueing information that simplifies the content of A.

(II) The student's response to frame P is taken as indicative of his psychological makeup, say as a student "type x" or a student of "type y". Suppose the student is "type x", he is subsequently assigned to a subprogramme X; if he is "type y" he is subsequently assigned to a subprogramme Y.

It is evident that writing a "Branching Programme" is a much more demanding process than writing a "Linear Programme" and that a comparatively elaborate teaching machine is required to administer a "Branching Programme" (it has to detect different types of response and to make the branches encoded by the programmer). However, for certain sorts of subject matter and for student populations showing marked individual differences, the "Branching Programme" is a far more powerful medium. Further, it is able to capture and retain the student's attention much more effectively than a linear programme.

5. Computer Assistance

The sequential teaching machine can be replaced by a digital computer equipped with a suitable display and response console (which usually consists in an automatically controlled slide projector and cathode ray tube for display and a teleprinter keyboard and a "light pen" for response). Many consoles and consequently, many students, can be handled by the same computer and the computer is also able to collate performance data and prepare statistics.

In the most straightforward case, the basic instructional programme is a branching programme of the type we have already considered. But the computer is well able (I) To provide more comprehensive "knowledge of results" than mere "all or none reinforcement signals" and (II) To accept more comprehensive responses than mere "button pressing" selections. The student may, for example, type in a whole statement or a whole equation. Further, (III), the branching in a computer controlled system can be determined by historical data gleaned from the student's past performance (rather than depending upon his immediate response) and this facility offers very significant advantages. Finally, (IV), the student may be allowed to interact with the machine in a completely different mode; he can be invited to make independent enquiries regarding the problems that are posed and to investigate the field on his own account. This mode of interaction (which could not be practicably instrumented on a simple teaching machine) entails the higher level discourse we cited earlier in this paper; it adds enormously to the value and efficiency of the teaching process and is probably an indispensable component of future systems.

6. Skill Teaching

Instruction of a skill is based upon a more or less complete simulation of the real life task (incorporated in a simulator for an aircraft or for air traffic control; in a real machine, such as a typewriter; in a set of test instruments, in the case of fault finding and diagnosis). The student is presented (through the simulator) with the real life problems posed by the skill.

It is rarely possible for the novice to tackle these real life problems as they stand. Consequently, specific training procedures are adopted. These are:

(1) Instead of a programme based upon distinct concepts (the expedient used in sequential teaching) the skill is reduced to component subskills, each of which demands the exercise or application of a concept. These subskills are separately rehearsed (as when a typist is trained using part of the alphabet only).

(2) The organization of the skill is built up by the joint rehearsal of subskills, due attention being paid to the effects of interference and the advantages to be gained from positive transfer of training between one subskill and the other.

(3) Knowledge of results is provided, at least in the initial phase of training.

(4) Cueing information is provided to simplify the problems posed by the skill.

All of these procedures can be regarded as modifications in the task simulator or its input.

In the present paper, I shall concentrate upon training systems in which these modifications are carried out, as a function of the student's performance, by an adaptive teaching machine.

7. Adaptive Machines

The adaptive machine continuously measures the student's performance in the simulated situation and forms statistical estimates of his proficiency. It is important to notice that these measures may be inherently complex. A student can make different classes of mistake and, in general, there are as many separate proficiency measures as there are distinct classes of mistake.

The "knowledge of results" signals delivered to the student by the adaptive machine are obtained from the same performance comparator (a device for comparing what the student does, with what he should have done) that feeds the measuring circuits, (however, the proficiency measures generally entail more complex properties than "right or wrong"; they depend, for example, upon how fast the student responds).

In terms of these measures, the adaptive machine builds up a description of the student's behaviour in its "memory" circuits. To give some idea of the amount of "memory" involved, we found it necessary, in designing a machine to instruct teleprinting, to use 160 separate analog memories arranged in a 5 by 32 pattern; there being 5 types of proficiency measure; 32 distinct types of stimulus.

The adaptive teaching machine can modify the training conditions by arranging different rehearsal sequences for the sub-skills and by providing cueing information that simplifies the problems that are posed.

In teleprinting, for example, the student is required to solve the problem posed by the appearance of a character in an alpha-numeric display (he solves it by selecting the right key on the teleprinter keyboard within an allowed interval). Sub-skills correspond to types of text. The cueing information consists in a specification (either visually or by auditory and tactile signals) of the sort of motion he should make. The <u>amount</u> of cueing information depends upon the delay in its presentation after the appearance of an alpha-numeric stimulus. There are as many <u>sorts</u> of cueing information as there are types of mistake.

The machine uses its behavioural description to determine a rehearsal sequence and the amount of each sort of cueing information that should be delivered to simplify the problems. For this purpose, it has an inbuilt teaching strategy, based upon a specific learning model and intended to satisfy the conditions and criteria outlined in the first part of this paper. A typical strategy presents the novice with a great deal of simplification. As he becomes proficient, this co-operative assistance is removed, but it is removed selectively (so that the student is still given assistance with respect to problems that he finds personally difficult). The process is also reversible (if a student gets into difficulties and begins to delay his response or to make mistakes, he is given co-operative assistance). Ultimately, when the student has learned the skill, he is required to solve real life problems without assistance, and, at this stage, he is, by definition of the system, capable of doing so.

Machines of this sort are able to establish rather a close rapport with the student and they have been successfully used for the instruction of many industrial, military and intellectual skills. The relation between the student and the machine is game-like and conversational. But strictly, the discourse is limited to a single level; a level of stimulus and response; or problem and solution. Because of this, there is a tendency for the student to engage in a game like interaction, whereby he tries to gain control over the teaching strategy of the machine. He does so by learning the characteristics of the machine and making responses that the machine will interpret as mistakes, even though the student is not, in fact, experiencing difficulties.

This mode of response is "illegal" within the tenets of the system. But, as we argued in the first part of the paper, it is to be expected. Hence, as teaching machine designers, we need to improve our machines.

8. Adaptive Metasystems

The initial refinement of the system consists in providing a channel for the sort of higher level discourse that occurs in a tutorial conversation when the student asks the instructor to rehearse different problems or teach in a different way. Once again, the statements are "mechanized" rather than verbal, but they are statements about the stimulus response interaction and are thus "metastatements".

The real-life instructor does not accept the student's preferences and proposals unconditionally. He examines them and is prone to accept them if, and only if, he believes that the strategy posited by the student would be more conducive to learning than his own choice of strategy. Hence, the next refinement in the system mimics the conditional acceptance of a real-life instructor. Broadly, we set up a mechanism whereby the student is allowed a variable degree of control over the teaching strategy. The student vies with the machine for this control (one or the other of them must determine what is done). Such a system (called an "adaptive metasystem") is designed to allow the student a degree of control that is proportional to the relative effectiveness of his own proposals. If, as a result of his previous proposals, the instructional goal is approximated, he is allowed more control at the next move. If the instructional goal is not approximated, he is allowed less.

This sort of rule leads to a state of affairs in which the novice has no control of the teaching strategy; the "partially proficient" student may have control; the completely proficient student loses control since he has ceased to learn and at this stage the basic machine strategy impels him to deal with real life problems in the absence of any assistance from the system.

The man-machine interaction in an adaptive metasystem is logically akin to a conversation. Indeed, it is probably fair to claim that in all relevant respects it is identical with a conversation. Certainly, the teaching efficiency of these systems is very high and there is little doubt that future designs should embody this paradigm.

In conclusion, we comment upon the similarity between the adaptive metasystem (for skill teaching) and the enquiry mode (in the sequential teaching machines). In fact, at this level of elaboration, the distinction between sequential and skill teaching loses its potency and the joint emphasis upon a full conversational interaction indicates, I believe, the trend of development in the field of cybernetically oriented education.

unclassified

Security Classification

DOCUMENT CONTROL DATA - R & D

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) System Research Ltd 20 Hill Rise Richmond, Surrey, England	2a. REPORT SECURITY CLASSIFICATION unclassified 2b. GROUP

3. REPORT TITLE

CYBERNETICS AND EDUCATION

4. DESCRIPTIVE NOTES (Type of report and inclusive dates)

scientific; ; interim

5. AUTHOR(S) (First name, middle initial, last name)

Gordon Pask

6. REPORT DATE	7a. TOTAL NO. OF PAGES	7b. NO. OF REFS
Aug. 1967	14	0

8a. CONTRACT OR GRANT NO. F61052-67-C-0010	9a. ORIGINATOR'S REPORT NUMBER(S)
b. PROJECT NO. 9769-04	
c. 61445014	9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)
d. 681304	AFOSR 67-2001

10. DISTRIBUTION STATEMENT

Distribution of this document is unlimited.

11. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY
	Air Force Office of Scientific Research Directorate of Information Sciences Arlington, Virginia 22209

13. ABSTRACT

This paper reviews the main contribution of cybernetics to the art and science of teaching. It reviews the bearing of cybernetic principles on teaching methods, and considers teaching as the control of learning. It describes the difference between sequential and skill teaching machines and the value of computer assistance in teaching machines. Adapative machines and adaptive metasystems are briefly discussed.

DD FORM 1 NOV 65 1473

Security Classification

Security Classification

14. KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Cybernetics						
Teaching						
Learning						
Teaching machines						
Adaptive systems						
Control systems						
Computer-assisted instruction						
Man-machine interface						

Security Classification