ADAPTIVE TEACHING MACHINES

By

I. LAKE

*Cybernetic Developments Ltd., London**

PRELIMINARY CONSIDERATIONS

NO training programme can be efficient unless it embodies procedures for resolving the student's difficulties. Skinner-type programmes try to overcome this problem by advancing the student in small easily-comprehensible stages that illustrate as many facets of the subject matter as possible. Since the student is tested at every stage by a question, guided by the representation of a correct answer, and prompted wherever necessary by hints etc., there should in fact be no difficulties that the student cannot readily *evaluate for himself*—e.g. by taking another look at the instructional material immediately preceding.

Skinnerian methods exhibit what might be called a "piecemeal" approach to understanding. Information is supplied in easily-digestible units so that the student can sort out his own mistakes as he goes along. Not every subject can be satisfactorily taught in this way. Whereas some subjects (like spelling) can be comprehended as a sequence of relatively simple and independent facts, other subjects must at some stage be grasped as a totality because they involve facts and rules that are intricately related to one another. If an elaborate network of relationships has to be built up it may be desirable to practise the student on several inter-related facts and rules at a time. This naturally makes the units of instruction less digestible and may well create difficulties that the student is unable to overcome by his own efforts. In such cases it is up to the teacher to provide more positive assistance, by evaluating the student's misunderstandings for him and by taking him through an appropriate correction procedure.

Machines like the Autotutor and Grundytutor (which embody the branching technique) are just about the simplest devices that can claim to provide this facility. After each unit of instruction, the student is shown a test question and a set of possible answers (marked A, B, C,.) of which only one answer is correct. The student must locate the correct answer by pressing a corresponding button (also labelled A, B, C,.) and if he chooses wrongly he is automatically switched to a "remedial lesson" appropriate to the kind of mistake made.

Thus if a student asserts that angle ⌋⌊ is greater than angle ⌋ an appropriate remedial lesson would explain that the size of an angle remains unaffected by variations in length etc. of the straight lines involved. Note that the question and answer technique is here being put to a different use. Whereas the Skinnerian method protects the student against error (by simplifying instruction to a point where errors are hardly possible), the Autotutor encourages the student to make errors that will reveal the nature of his misunderstandings.

Of necessity, the Autotutor operates with somewhat limited evidence. On the strength of a single question, the student is categorized as being right or as being wrong in just one of a limited number (usually 2 or 3) of ways. Since there is a pre-set correction procedure for every wrong button pressed, it is tacitly assumed that everyone pressing the same wrong button is guilty of much the same misunderstanding.

This assumption is bound to break down at times.

And when it does the student receives a remedial lesson that partly or wholly misses the point because it does not impinge directly upon his misunderstanding. This may well confuse and irritate the student. But if he perseveres he will usually come across further instructional materials that are sufficiently "remedial" to set him back on the right lines. If the student is motivated to overcome such irritations (which may be quite small compared with the irritations that beset conventional teaching), the whole procedure is likely to be very successful indeed. In fact it is not at all uncommon to hear of students learning in less than half the usual time.

Although the method has achieved some dramatic results, it undoubtedly has its limitations. As already indicated, each remedial lesson is based solely on the student's last mistake; and it is often very difficult to gain reliable information about the student in this way. At one extreme, a student might press the wrong button because of some fleeting mental aberration. Or he might do

*It had been arranged for Mr. Lake to come to South Africa to take part in the Conference, but owing to pressure of work he had to decline at the last minute. Instead he sent this paper to be read at the Conference.

so because of some deeply-rooted misunderstanding. To be on the safe side the remedial lesson must therefore be sufficiently laboured to cover the latter case, and it must also examine the mistake from as many different viewpoints as there are (likely to be) misunderstandings. Yet even this does not get to the heart of the matter. For it remains true that single questions are rarely sufficient to locate the deeper sources of the student's difficulty. In consequence, the remedial lessons that the student receives can often do little more than nibble at different aspects of several (possibly unrelated) misunderstandings. This may well give the students a *fragmentary* understanding of the subject matter. If he gets anything more out of the instructional material it is because he is able (or, more pertinently, *willing*) to do the machine's job by synthesising the fragments for himself.

The fact that so many students tend (despite manifold frustrations) to make the necessary syntheses, says a great deal for their perseverance and their attitude toward the educational process. It also enables such methods to claim successes that they would not otherwise have achieved. But some students will invariably fall by the wayside, especially if the synthesising process presents special difficulties. For these stragglers a more businesslike procedure is required. And it is here that the adaptive teaching machine has an important part to play.

Its general properties.

AN adaptive teaching machine is a device which in a non-trivial sense *learns* about the student in order to teach him more effectively. Unlike the simpler machines, it can store information (on internal registers or "memory units") about the way in which the student has handled the various test questions given to him. It can examine a whole sequence of mistakes to see whether they are collectively indicative of some fundamental misunderstanding. And it can manipulate correctional procedures in a much more indicative manner.

Like the Autotutor, the adaptive machine provides a sequence of instructional items on some suitably automated display, and a set of multiple-choice test items for determining (at frequent intervals) whether or not the instruction has been understood. But it deals with the student's answers (to these test questions) in a highly distinctive way. In fact it uses them to construct, on internal registers, an up-to-the-moment account of the student's relative strengths and weaknesses. The information accumulating on these registers is then fed through to a small computing mechanism which decides what changes, if any, should be made in the training routine.

This is less complicated than it sounds. The internal registers or "memory units" need consist of nothing more elaborate than a set of reversible counters which go up when the student answers his test questions correctly, and down when he does not. And the computing mechanism might simply be a set of relays that remain inactive until one of the counters clicks round to some pre-set critical value.

In this particular case, the attainment of some critically high value might cause the machine to switch to a more challenging set of instructional materials. And the attainment of some critically low value might cause the machine to present (on a supplementary display, for instance) a set of hints and cues which are likely to help the student out of difficulty. The machine therefore has the capacity to adapt in two different ways. If the student is doing well it presents harder problems to spur him on. If the student is doing badly it presents easier problems to help him out.

This simple arrangement is fairly typical, and can be elaborated in many different ways. Variations in the states of counters could, for instance, be made to vary the time allowances available to the student for answering different kinds of test questions. Whatever the degree of elaboration, the basic picture remains the same. *In every case some cumulative record is kept of the student's efficiency, and this is used as a basis for adjusting his training programme in a sensible way.*

To be efficient, the machine must obviously maintain the right kind of cumulative record. And it turns out that the best procedure is to evaluate on the machine's registers the student's ability to surmount those obstacles that most commonly frustrate the learning process. Since this point is of central importance, and is not at all self-evident, the next section will consider it in greater detail.

Mode of approach

In their efforts to teach efficiently, teachers often break down their lessons into what they take to be a "natural cluster" of (say) three or four "easily-relatable" ideas or concepts. In such cases instruction is generally alternated with questions aimed at testing the student's ability to handle such concepts—both singly and in relevant combinations. If a question is answered wrongly, the student is typically given a verbal exposition relating certain details of the question to the concept(s) it was meant to exemplify. The original concepts are then reiterated and exemplified in further exercise materials until the teacher is finally satisfied (perhaps wrongly) that the student has got the message.

This can be a highly inefficient and exhausting procedure. The trouble is that the student, who does not see things in the same way, may well

suffer from misconceptions that cut right across the teacher's concepts and questions. There will therefore be occasions on which the teacher imagines that the student has seen the point. And there will be occasions when his hopes are dashed. By exemplifying umpteen different facets of the concepts he is trying to get across, the teacher may eventually bulldoze the student into giving (grudgingly, perhaps and without conviction) the answers required of him. But instead of hammering away at the ideas which are clear only to the teacher, it is much more to the point to make a direct attack on the student's misconception—to root them out and eliminate them so that the student can gain a genuine understanding of the subject matter being imparted.

These considerations imply two fundamentally different approaches to machine instruction. One of these adopts a therapeutic approach to the student's difficulties. The other seeks to suppress such difficulties by imposing a different point of view. The latter approach is essentially a form of indoctrination, and machines built to operate in this way need only have a separate register for each important idea (or combination of ideas) being pushed. These registers would advise the machine which aspects of the indoctrination had been assimilated, and which had not. The machine would then simply load the student with more and more exercise materials etc., on those aspects for which the registers continued to denote inadequate assimilation. In due course the student might come to understand the ideas in question, or he *might* just learn how to convince the machine that he understands.

The alternative procedure (and the procedure used at Richmond) is to design the instructional and exercise materials so that they show as much sensitivity as possible to the student's difficulties. In this case prior research must be carried out, over a large population of students, to determine exactly what these difficulties are. Usually it will be found that the students make errors, many of which really do fall into "natural clusters"—because they are in fact occasioned by the same underlying misconception. Given this knowledge, the machine can then be made to organise its exercise materials into distinct classes or *problem areas* which impinge rather directly upon the classes of misconception that are known, from actual research, to stand in the way of learning. When a student makes an error it can therefore be categorised as reflecting, in all probability, a particular kind of misconception. And when enough evidence has been accumulated to suggest that the misconception really is present, the machine can be made to apply a remedial procedure that is likely to be genuinely apt and beneficial because it is conducted at the right depth.

Under this procedure, the machine's internal registers must be put to the basically *different* use of detecting the nature and extent of any misconceptions that are present. And the machine must also have facilities for acting on the information that accumulates in these registers. In order to see how this is done it is necessary, however, to say more about the way in which the machine works.

Internal organisation

The machine's internal organisation can be broadly specified as follows:

1. At least one register must be provided in respect of every misconception for which evidence needs to be accumulated. If two or more misconceptions are likely to interact in ways that cause special difficulties, separate registers may also be necessary to detect the presence of such interactions.

2. Each test item must be associated, in a quantifiable way, with one or more of these registers. In other words each test item must contribute at least some information about the student's ability to overcome standard difficulties. If the instructional material involves mastery over (say) three or four distinct misconceptions or problem areas, the machine will therefore contain at least three or four separate registers and the test items will be correspondingly grouped into at least three or four separate classes. A test item intended to test the student's mastery over problem area 2 would thus be coded (for the machine's benefit) to ensure that only register 2 is activated by the student's response.

3. If necessary, test items can be designed to test the student on several problem areas at once. If a particular test item is likely to involve misconceptions 1 and 3, a correct (or incorrect) response could readily be coded to produce a uniform increment (or decrement) in each of the registers 1 and 3. And if the test item involves misconception 1 to a greater extent than misconceptions 3, the code could be further adjusted to produce differential increments (or decrements) in these two registers. Presumably it is obvious that prior research must be carried out to establish how these test items should be coded. Given a judicious coding system, the registers can rapidly provide useful cumulative indices of the student's ability to handle different aspects of the instruction.

4. To complete the system, the machine must embody a mechanism which can evaluate the information accumulating on these registers and can cause the training programme to vary in ways that seem most likely to meet the student's special requirements. This mechanism has been described, in the teaching machine literature, as an *overall control mechanism*. Like a real-life supervisor, it examines all the registers that are active in order to see, for each in turn, the extent to which the student is falling short of some ideal criterion of efficiency.
   In the simplest case, this mechanism is a simple comparator which matches the state of each register against some pre-determined norm and secures an electrical "read-out" of the discrepancy. The greater the discrepancy, the greater is the student's need for assistance.

5. By coupling the electrical read-out to a supplementary display on which hints and remedial exercises etc., can be given, the machine can provide a greater or less amount of help according to the magnitude of the discrepancy in question. This coupling is, however, seldom direct. For in some cases the registers will show the student as requiring more help than the machine can give at any one time. The overall control mechanism must therefore *integrate* the information it receives from different registers in order to determine the best compromise procedure.
   The kinds of computation needed to establish this compromise are usually determined by decision rules that are programmed into the machine at the time of its construction. But there could very well be a facility for changing these rules in the light of experience. And the machine could even be programmed to change its own rules if it found that certain procedures were not working effectively.

6. The adjustments that the machine makes to its training routine are not *exclusively* concerned with resolving the student's misconception. In every dynamic training process, the need to help the student when he is doing badly must always be weighed against the desirability of spurring him on when he is doing well. If the machine's registers show that the student has not reached an acceptable standard on certain problem areas, he must continue to be given supplementary assistance on these areas. But once the required standard is reached, it is equally important that the machine should switch the student to a harder batch of materials. This means that the instructional and test items associated with each problem area should be internally grouped into sets of increasing order of difficulty—so that the student can be rapidly shifted onto a set appropriate to his level of ability.

7. The operations outlined above can be carried out with varying degrees of sophistication. Since each new response by the student alters the state of at least one register, the record of his performance is being repeatedly up-dated. In principle, this could lead to repeated adjustments to the training routine. But it often suffices to defer adjustment until one of the registers reaches some critically high or critically low level.

   The registers might, for instance, simply count (for different problem areas) the excess of correct responses. In this case the machine might wait until this excess figure reaches + 8 (say) on any register, and then switch the student to a more advanced set of exercise materials on that particular problem area. If the counter then slips back to below + 8 the student is duly returned to the more elementary level. But if it reaches the next critical level of + 16 (say), the machine promptly reduces the number of exercises given in that area. This enables the student to devote more time to other problem areas on which he has not yet made the grade. While the counters are moving backwards or forwards towards their critical values, the machine could be made to adjust certain *time* variables—such as the time allowed to respond to certain kinds of test items. In general this would mean that if the student gives a correct response to a question involving (say) problem area 1, then the next question involving that problem area would have a slightly shorter time allowance assigned to it. Conversely, a mistaken response would produce a slightly longer time allowance on the next occasion. Since humans are on the whole rather sensitive to such time adjustments, their manipulation greatly heightens the student's impression that he is receiving a personal and "tailor-made" training routine.

   In cases such as this, the student is entirely right to regard his programme as tailor-made. Whereas the Autotutor sends everyone along the same main stream of instructional items (admittedly skipping and branching here and there), the adaptive teaching machine organises its instructional materials into groups which can be presented *in any order* and *at any pace* that the student seems to require. There is no "main line" that is common to all students and the programme is, in a very real sense, *constructed* by the machine as it goes along.

   Furthermore, the making of one kind of adjustment in no way prevents the machine

from preparing another. If a student has just made one mistake that points to misconception 1 or 2 or 4 and another mistake that points to misconception 2 or 5, the two mistakes *collectively* suggest that misconception 2 is predominant. If the machine now switches to a remedial procedure for the elimination of misconception 2, this does not preclude the possibility of continuing to secure more evidence about other possible misunderstandings.

8. The technically minded reader might like to note that the adaptive teaching machine has the logical calibre of an *hierarchial control mechanism.* To be non-trivially adaptive, the machine must have at least two registers accumulating information. And it must have an overall controller that scans the states of these registers, performs certain predetermined computations and adjusts its training programme according to built-in decision rules.

   This is the simplest (two-level) form of hierarchy. In more complicated cases there might be two independent sets of registers (one recording speed and the other recording accuracy) which are cross-coupled or related at some higher level. Alternatively there might be some intermediate level of registers concerned with the *relationships* between pairs of registers at the lowest level.

   It might happen, for instance, that the overcoming of one misconception leads to the development of another. If there is this kind of interaction or "negative transfer" among different problem areas, the state of one register might be inversely correlated with the state of another; and it requires at least one higher-level register to detect the presence of this correlation. These are technicalities that cannot be gone into here, but it is worth remarking that a precise understanding of such phenomena is contributing a great deal to the public stock of knowledge about the teaching process.

General teaching strategy

The adjustments that the adaptive teaching machine makes to its training programme will vary according to the complexity of the machine and the nature of the subject matter being taught. Some adjustment procedures will necessarily be more complicated than others, but all of them are subject to the vital requirements that they should keep the student *motivated.* Unless the adjustments have the right sort of impact, so that they continue to enlist the student's interest and attention, the whole procedure is likely to fall to the ground.

As a first step towards keeping the student motivated, it is essential to adjust the training programme to a level that the student is likely to find congenial. If the instructional materials are too simple, the student will get bored and his attention will wander. If the instructional materials are too difficult, he will get frustrated and (once again) his attention and enthusiasm will diminish. By moving into a region which is neither too simple nor too difficult the machine can, however, give the student problems which he regards as challenging and (more importantly) as being within his immediate realm of competence.

To locate this critical region, and to stay within it if it shifts, the machine must adjust the standard of difficulty of its instructional items in a manner determined by the student's own performance.

And this can be done by reference to its internal registers. If these registers show that the student is making very few errors, the difficulty is stepped up until the error score starts to rise. At this point the machine eases off until the error score starts to fall again, whereupon the pressure is gently reapplied.

By shifting its level of difficulty in this way, the machine guarantees that the student will receive a regular supply of challenging problems. In this respect the machine satisfies one of the most fundamental human needs—the need for new and stimulating experiences. Looked at another way, it can be said that the machine adopts a pressuring technique which endeavours to keep the student working as near as possible to the threshhold of his ability. This is an ideal state of affairs because it means that the student cannot afford to let his attention wander far from the job. In consequence, he tends to be optimally amenable to the teaching process.

As a bonus, it turns out that when the student is working near his peak he may well reveal incipient weaknesses and misconceptions that might otherwise not have shown up. So the pressuring technique also tends to maximise the amount of useful information that the machine can gain about the student. (To prevent the student from becoming embarrassed by the weaknesses he is revealing, the machine can often conceal the fact that it has noticed such weaknesses. This protects the student's ego, and enables the machine to go on collecting information without losing the student's goodwill.)

Perhaps it should be stressed that although the machine may succeed in temporarily pushing the student to the peak of his ability, it is up to the student to decide whether or not he wants to stay there. If he suddenly tires of working under pressure he will inevitably start to make mistakes

(e.g. by failing to answer a test item quickly enough) and the machine will promptly slacken its pace. As might be expected, this whole procedure tends to generate an extremely congenial relationship between the student and the machine —not unlike the relationship between a child and a novelty-producing toy. Once the student starts interacting with the machine, he tends to get permanently caught up in a stream of activity. And this is hardly surprising because the machine is deliberately engaging him in a sort of intellectual donkey and carrot race. In the early stages of instruction he is offered inducements (in the form of hints and similar guidance) that enable him to improve, but as he gets better the inducements are systematically withdrawn.

This situation can be summarised more formally by saying the machine engages the student in a partly co-operative/partly competitive game. By giving the student special hints and guidance the machine clearly does some of the work that the student's brain is not yet capable of doing. In this respect, the machine can be said to function as a co-operative extention of the student's brain. But the machine is a reluctant co-operator that refuses to spoonfeed the student. It repeatedly tries to withdraw its assistance, in a selective manner, so that the student can stand on his own feet. And this is the competitive element in the game. It arises because both the student and the machine have the same ultimate objective—namely, to reach a state of affairs in which the student is able to handle all the most difficult problems with the minimum of support.

The adaptive teaching machine can therefore be described as a device which forms a close and dynamic partnership with its student in order to maximise some quantifiable measure of the student's efficiency. To this end, the machine sometimes co-operates with the student by helping him out of difficulty. And it sometimes competes with the student by putting new difficulties in his way. But both these strategies have the common aim of keeping the student on the move. By withdrawing its assistance at regular intervals the machine finds out what the student is worth, and endeavours to keep him working at high capacity. In a very real sense, the machine minimises its own efforts so that the student must maximise his. But active support is always ready to hand whenever the student shows signs of needing it.

Given that the student is motivated to learn, effective teaching is largely a matter of imparting knowledge in a way that seems most likely to resolve the student's misconceptions. It is the claim of this paper that adaptive teaching machines have facilities for keeping the student motivated and resolving his difficulties, that the simpler machines do not possess to anything like the same extent. They should therefore be superior to the simpler machines, but they are of course considerably more expensive.

In all probability there are many school subjects for which it would be both uneconomic *and* unnecessary (especially in the case of "normal" students) to use the adaptive teaching machine. And some non-partisan research is currently being conducted at Richmond to ascertain what kinds of subjects can most profitably be taught by what kinds of methods.

Undoubtedly, there are profound differences between the adaptive teaching machine and (say) the Skinner-type machine. And this leads one to suspect that they might best be used for different purposes. Whereas adaptive machines keep the student working at the highest capacity congenial to him, the Skinnerian machines tend to keep the student working at the lowest capacity congenial to him. They therefore operate at opposite ends of that narrow range of problems which are not-too-easy-but-not-too-difficult. By ensuring that students give correct answers to almost all the test items the Skinnerian machine necessarily presents problems that have only minimal intellectual bite. So the student is working in a context in which he is seldom fully extended.

This may be fascinatingly hypnotic, especially to persons who expect the educational process to be somewhat painful. And it may well be an optimum technique for conditioning people to recite the Capitals of Europe or to perform some similar "rote memory" task. But there are reasons for doubting the long-term wisdom of teaching in this way. In the first place it equips the student rather poorly for the outside world, where learning usually is painful. Secondly, it is probably inadequate for the teaching of any form of *skill* in which there is some rather small body of basic ideas or axioms etc. (as in mathematics, logic, science, jurisprudence, political philosophy and the like) which the student must learn to manipulate in an efficient way over a wide and unpredictable variety of circumstances.

On the other hand, the adaptive teaching machines also have their defects. For instance, there is a certain kind of student who becomes acutely anxious if time limits are imposed on his behaviour. In principle, the adaptive machine could be made to test for this kind of aversion. Alternatively, the student might be provided with a knob which enables him to remove the time constraints under certain conditions. But all this adds both to the elaboration and to the cost of the machine.

Generally speaking, special-purpose gadgets of this kind are subject to the law of diminishing returns. It is not possible to cater for every human quirk, and there is little point in bothering about the odd one or two. Perhaps the best procedure is simply to equip the machine with a *distress signal* which is actuated whenever the student-machine interaction reaches an impasse. (Whenever, for instance, the machine exhausts some set of exercise materials before getting the student to the standard demanded by the follow-up set). The distress signal would of course summon a real-life teacher for specialist and individual guidance. And this is exactly the function for which the human teacher is best suited. Moreover, it is a function that he would be able to perform much more efficiently because the machine would provide him with a rather specific account (on punched tape, if necessary) of what had gone wrong.