AUTOMATED TEACHING
The Problem and the Prospects

By

RICHARD GOODMAN

Head of the Department of Computing, Cybernetics and Management, Brighton College of Technology

AS with every new technological development, there has been and is an underdamped upswing of interest, as Huskey has put it, in the case of programmed instruction. There comes a stage when it is advisable to exercise a little judicious damping so that we can pause for a moment and ask ourselves what in fact we really are doing and what we really are talking about. For there is a very great danger that we are using words rather facilely, without in fact, really understanding what we are using them to say. So one of the things I would like to do is to ask some questions which may help us to consider what we are in fact talking about and what we are in fact after.

Let us begin by stressing the very great incongruity existing at the present day, an incongruity which is yet very often passed by almost without comment. We live in a world where, technologically at least, we have the means, or shall have the means within a very few years, to tackle, and be well on the way to solving, most of the major problems which confront humanity—hunger, disease, unemployment, ignorance. We have the technological means to do this *and* we have the technological means also to begin our voyage to the stars. We have the means, but we lack the skilled people without which we cannot get down seriously to tackle these problems. This is true not only on a world scale, but it is true in individual countries. In my brief stay here in Africa one of the most exciting things, I think, is to feel that here we have a world which reflects as it were—it's hardly right to say in miniature—the whole world in which we live—a very exciting world with fantastic promise. The key to the fulfilment of that promise is the development not only of the skilled people required to carry out the jobs—but also the creation of a background of education and imagination among the general public to enable them gladly to support these enterprises and to understand them. And yet our educational technology, upon which the supply of this imagination and these personnel depend, is many hundreds of years old. We are still in the days of Peter Abelard and of Caxton. Our educational technology is hundreds of years out of date. It is also true, of course, that many of the subjects we are teaching today and much of the whole organisation and of the content of education is out of date. Daily in the schools, murder is committed. I hesitate to think of the number of promising mathematicians who have been killed as mathematicians because they were taught arithmetic in the primary school. But this, the content of education, is not really our subject for this morning, although I do want to emphasise that the introduction of new technological devices, the design of a new educational technology cannot go forward without, at the same time, a considerable investigation into what it is that we are teaching and how we are teaching it.

How then can we go about modernising and re-designing our educational technology? And why should we do this? Expediency demands it, but it seems to me that any argument for expediency is not sufficient. We have to ask ourselves both whether it is possible in fact to automate part, at least, of the learning-teaching process *and* whether it is desirable to do so. I think the answer to the first question will indicate a positive answer to the second. In undertaking this investigation we need also to remind ourselves that no means of communicating information is sacrosanct.

The lecture type of educational technique is, for instance, a very inefficient method. We know that in a situation such as this in which I am speaking today, seventy-five per cent of the information imparted will be lost after about eight hours. You may remember some twenty-five per cent of what I am saying for many years to come. But it is usually the most irrelevant and unimportant parts of what is said. The printed book has been a very fine way of imparting information for a very large number of years. But why, as our technology develops, should we hanker to preserve a method of imparting information which was suited to the days of Caxton, and perhaps Oxford of the last century, but is probably not today? We should not regard the book as such as sacrosanct.

Where, then do we turn? I suppose it would be natural to turn in the first instance, because learning at least is supposed to be a psychological process, to the psychologists. But, alas, what do we find? Professor Meredith of Leeds University estimated a couple of years ago that there were 29 different theories of learning. There is in fact no generally accepted, working psychological theory of learning. This is unfortunate, but I think it is almost inevitable. Probably one of the best ways of developing such a theory is to use and experiment with the kind of devices we are here to talk about today. If we look at most psychologies today, what do we find? We find that psychology, at least theoretically, is still very heavily tied to the matrix of thinking from which it has still not completely separated, the theory of metaphysics. We find, for example that most psychological theories have a great familiarity, a great similarity, to a monadology of Leibniz, which was an extremely atomic sort of philosophy leading to solipsism. And in fact most psychological theory today is a form of dishonest solipsism. Since solipsism is, apart from anything else, an extremely impolite philosophy, it is hardly one that appeals to those of us who live in a very real, integrated, dynamic society. If you examine most psychological theories, from pure behaviours to the more sentimental kind of romantic individualism that is occasionally served up as psychological theory, you will find that ultimately all of them, almost without exception, reduce to a form of solipsism. Although we have a mass of experimental evidence about the way people learn, we certainly have no unified generally acceptable theory. So where do we turn?

I think we must adopt another approach altogether. Instead of adopting the molecular kind of approach which starts with little bits and starts to build up, we have to take a molar approach, an over-all approach. Perhaps I can express the difference between these two approaches by asking you to consider the problem of synthesizing a chair, with molecules. If you were to get even a quarter of an inch up one of the legs, by that time you would be in such a state of indeterminacy that you would give up in complete and utter disgust. It would be better to employ a carpenter!

No, the approach indicated would seem to be somewhere along the following lines. Learning and teaching and human intercourse are communication problems. They are problems in communication and control just as well as any other natural process. It would seem then that although we don't know precisely how learning takes place, at least we can say *when* learning occurs. May we not then apply the ideas of communication and control techniques to the problem of human learning? We can do this today because we have built up in recent years a considerable amount of knowledge and know-how in these matters. We know a great deal about data processing as a result of our experience with high-speed computers. We have a fairly developed information theory behind us. We know a fair amount about self-organising systems, about coding and what this involves. I would suggest therefore that this is the approach for the moment. Ultimately a good, sound and generally acceptable theory of learning may be developed as a result of these efforts, but even then it will probably not be as "psychological" as some anticipate.

If, then, we adopt this approach, where do we begin? Consider a simple organism in its environment, being careful not to fall into the fallacy that Whitehead pointed out of *simple location:* because the learning organism, the organism in its environment is essentially co-extensive with that environment. Learning (and no organisation can exist without learning) is basically a question of the interaction of the organism with its environment. In lower organisms it is adaptation *to* the environment, but with higher organisms, and especially higher social organisms, the organism begins to learn also by adapting the environment to itself or to the group. Now when this happens, and for this to happen, several interesting things occur. Firstly, because the environment is infinitely complex, it is necessary to select from it and build up a model of that environment which is *homomorphic* with the environment itself, which behaves in some essentials as the environment behaves; we build up a set of synthetic, simplified, dynamic, and to some extent controlled, artificial environments upon which we operate in order to operate on the real world. We build a model, or a set of models for different purposes. Now if you are a social animal, then your knowledge of things is not only a knowledge by *acquaintance*, but a knowledge also by *description*. So you need in your model or artificial environments quite a lot of symbolism, of verbal presentation of information. The process of learning then will certainly not only involve a considerable amount of information processing, information retrieval, but will be concerned with symbolic materials, not necessarily verbal, but certainly symbolic.

What is then the role of the teacher in all this? Nobody in the long run can force anybody to learn anything at all. You can try to teach until you are blue in the face, but unless the desire to learn is there, learning will not take place. So we have to assume some primary motivation. This primary motivation, I believe, comes in the long run from the society in which you live. Its maintenance and development, however, is directly

part of the educational process. Is not then the role of the teacher precisely to create the set of simplified, controlled, dynamic environments we have been talking about and to guide as unobtrusively as possible the organism he is supposed to teach through them? We learn by being unobtrusively, but sometimes firmly, guided through such environments, in much the same way as we learn when we are guided in our exploration of a strange country. The accumulated experience of human teaching through the years confirms this. In the first place, effective learning takes place if the learning organism is encouraged to take responsibility for his own learning. In the second place, effective learning also takes place when the learning organism is allowed to proceed, in his exploration of the strange country, at his own pace, and that, when he is so allowed, he proceeds at a very much greater pace than we expect. In the third place the material we learn, or the journey we take through the material is best and most effectively made if we do it in smallish steps; in smallish steps after each one of which we are either reassured that we are in the right place, that we have done the right thing, that we have got where we wanted to go, or action is taken *immediately* to correct us if we have gone a little astray. (You notice I say smallish steps: there is an awful lot of nonsense talked, if I may say so, about "step-size". It is highly probable that the step-size depends on many other variables, and is not an independent variable in itself. And you can, of course, make your steps so small that you don't learn at all.) We also know that there are phases in the learning process when the material should be, as it were, "co-operative" with us, and there are other phases where the material must *challenge* us. There are phases in the learning process where the organism must be kept at stretch. During such phases you have seriously to take into account the probabilities of error. The idea that people learn best by not making mistakes may be true in certain sectors of the learning process. I feel, however, that it is not true when you come to problem solving-like learning. Consequently those programming structures which are based on partial psychological theories of learning which do not admit the use, if you like, of an error in the learning process need some kind of reformulation.

We must now ask the next question: can the teacher adequately fulfil the particular function we have ascribed to him? We said, if you remember, that in the process of guiding the organism through the sets of simplified, artificial environments, much information processing, information retrieval, symbol manipulation is involved. It is also true that, because we are social animals, and because even teachers are social animals, this environment must inevitably be shot through with values. May I put it in a slightly different way? In the learning-teaching process, there are two rough divisions of function. There is that function where we build models, perhaps for the sake of building, as we learn mathematics for the excitement and sake of learning mathematics. We build models and we set them to work just as we set to work a bit of Meccano and delight in seeing how it works. But there is the other phase where the models that we have set up and the output they have yielded must be evaluated in the light of the human experience of the world in which we live. I would like to suggest that the human being, the human teacher alone, can perform the second function which involves social evaluation, moral evaluation, care of the individual, acting if you like as catalyst of the imagination, maintaining and encouraging motivation. But whereas the human being can do this and no one else can do it, and in that function he fulfils the role of an educator, so far as information-processing is concerned, so far as symbol manipulation is involved, he is not a very efficient operator, and this for a number of reasons. First of all, he needs, in computer jargon, a lot of "down-time" and maintenance. Secondly, his very human, all too human performance, is variable and he is liable to become involved, engagé, if you like, and this distorts the information that he is attempting to transmit. Thirdly, precisely because his storage mechanism, his brain capacity, is so enormous, he finds it difficult to retrieve information easily and efficiently at just the right moment; and he finds it extremely difficult to reconstruct a logical process, once established, in order to arrive at the same result a second time. Fourthly, the very vehicle he uses to transmit information is what information theorists would call a very "noisy channel". Our language developed not really to state the case, but to express our reactions to what we take to be the case; therefore in using this language to transmit information all sorts of noise is inevitably introduced in the channel and the problem of teaching becomes, you see, not that of the noisy pupil, but that of the noisy teacher.

Should we not, therefore, allow the machine to take over the information-transmission, information-processing, the symbol-manipulation functions of the human teacher? Simultaneously thereby we shall also free him to become what he has never been yet; an educator, a guide, a person who is concerned primarily with the humanity of the student rather than with the symbolic and factual material that has to be transmitted.

How then do we go about automating the process, the function of information transmission? Let's again draw some distorted pictures. This

time I am going to depict an artificial environment by a rectangle, knowing full well that you realise that this is a monstrous distortion, and, once again, let's depict the learner by a centrally located circle. Interaction takes place between environment and learner. We shall put in—and there is no significance in the fact that I depict it as a square—the teacher. The teacher has, either externally or within himself, an information store, on which he draws from time to time. He has a selector mechanism and he will have a series of teaching strategies. Let us assume for our present purpose that he has only three such strategies. Some quantum of information is presented to the student. Within the strategy, the correct response is shown. The student responds and his response is fed into what one might call a *comparator*, into which the correct response is also fed. The result of this comparison is fed back to the teacher in two possible ways: either so that the *tactical sequence* of the teacher within a given stratagem is modified, or the current strategy itself is abandoned and another adopted. The important part in this diagram is this *feed-back* loop from *comparator* to *teacher* which enables the teacher to modify his tactics or his strategy according to the behaviour of the student. There are, unfortunately, many people who forget, when discussing programmed learning, that this feed-back loop from student to teacher is of equal importance to the feed-back from teacher to student, and that the feed from teacher to student is dependent upon the feed-back from student to teacher which modifies the material presented to the student according to the student's response. This closed-loop system is one of the most important things to remember when talking about programmed learning. Now we have to ask ourselves some more questions. Crowder, if you remember, maintains that with a teaching machine we are automating the "classical process of individual tuition". Other people have said that we are indulging in nothing more than the so-called "Socratic method". Others have said that we are doing nothing more than present the material in "logical form". I would suggest that all these suggestions are highly suspect and open to a great deal of questioning.

Are we in fact attempting to automate the "classical process of individual tuition"? This process is one which was developed in a relatively atomistic society. But today we are not living in such an atomistic society, much as some people might wish we were. We are living today in a society where, more truly than ever before, we are members one of another; moreover we know that much of human learning and much of human discovery takes place, and must take place, not with the individual alone, but as a result of inter-group interaction of several individuals. The individual cannot learn things unless he is a member of a group. So it seems to me that we are not in fact trying to automate the classical process of individual tuition. This is not what we are trying to do. We are trying to automate a certain phase, certain aspects, certain parts of the relationship between teacher, artificial environment and learning organism.

Are we doing this by a Socratic method? When I hear this word used, I get a little worried because I remember the impression which the Meno made on me when I first read it. You remember Socrates there was trying to teach a slave boy the Golden Rule, and he said something to him and asked him a question: Do you understand what I am saying? and the slave boy said (his constructed response was): "Yes, master"; then Socrates would say a little more and again ask: "Do you understand what I am saying?" and again the slave boy would answer "Yes, master"; by the time you have read the Meno all the way through you are quite convinced that Socrates was a complete humbug: all he was doing was to advertise his own knowledge; the slave boy had in fact learned nothing at all. It is perfectly possible to learn an awful lot about philosophy by reading the Platonic Dialogues, but it is not necessarily Plato's philosophy that you learn about. In short, it seems to me that to use the term "Socratic Method" is at least inexact; we are again using a terminology and an analogy which is not appropriate in a language which we must try to develop today to describe what we are doing. The meta-language we use to talk about what we are doing when we talk about teaching machines and programmed language can't be language built on analogy, it must be a more scientific language. Again, we are frequently told that when we programme, we must break the material down into a "logical" structure. I'm afraid that if you ever try to do this and then test your programme out in the context in which you wish to use it you are likely to find that the so-called logic of the material is vastly different from that of the learners you are trying to teach. In point of fact what we are trying to do is develop a form of organic sequencing. If we take the discovery approach seriously, the method of discovery that we adopt will depend largely on the people we are leading through this strange country. There will be, as it were, different "logics' for different types of people and for different subject matters. It is so very easy to say we split the material down into its "logical order" and leave it at that. I suggest that here again we ought to look very carefully at what we are saying.

I would now like to deal with two other points before we move on to other problems about the relationship of the teacher and about the structure of education generally. It has been said, perhaps rather exaggeratedly, but nevertheless with a certain element of truth, that, in the conventional situation (and the Rector has mentioned the prevalence of the word "fail" which should be banished from our vocabulary) the real function of education at the present day is to humiliate the student. When you think about it, this may be rather venomous, but there is a great deal of truth in it. If you recall, I said that in the Meno at least, it seemed that Socrates' main idea was not to cause the slave boy to learn, but to boost his own reputation, to show off. In point of fact, much of our teaching today suffers from the fact, not only that it is designed to humiliate the student, but is designed to make exhibitionists of the teachers. The teacher intrudes too much in the learning-teaching situation. Instead of being the unobtrusive reliable guide, he takes on a Socratic mystique, which he doesn't deserve, of "the teacher". Probably one of the greatest contributions which programmed learning and teaching machines will make to our educational system is to strip from the human teacher that mystique behind which he can dissemble and evade his duties. I am not saying that this is the responsibility of individual teachers, it is the result of the way in which our educational system has developed.

The information we impart to others is necessarily in verbal or symbolic form. There is a high percentage of verbal or symbolic material in the set artificial environments built in the course of the learning-teaching process. Can we also use teaching machines, programmed instruction with pre-literate training? The problem of *pre-literate programming* is here and elsewhere, of immense importance and urgency; yet very little has been done. It would seem, therefore, that, not only if we are to understand what we do if we are teaching at verbal level, but, for the other tasks which confront us, a concerted attack on the problems of pre-literate programming is essential. In this connection may I, rather arrogantly, as from the outside, pass on two bits of warning, because we ourselves have already been bitten. In the first place, it is, we believe, quite unimaginative to import programmes from abroad and expect them to suit our needs. If a programme offends in the tiniest detail against the *mores* of the society in which the learning-teaching process is taking place, you will undo every good that you could possibly have done, and the idea of programmed instruction will suffer a severe set-back. Those of us who don't know anything about Africa, but know a little about programming will always be only too pleased to help, but the ultimate responsibility for the development of this kind of teaching technique must lie with the people in the country in which the programmes are to be used.

The second thing I would like to say is that not only must we be careful not to offend against the *mores* of the society in which we are using this programming, but we have, also, to investigate the reaction of developing peoples to the whole concept of automation itself. What is going to be the effect of introducing teaching machines in a society of such-and-such level of development?

This is another question: Do we scrap then all our partial "psychological" theories of learning, all 29 of them? This would be the last thing to do. On the contrary, we must scrutinise them very carefully from the point of view that I have been trying to develop. And we use what is valid in them, together with the data that is beginning to be built up from the use of teaching machines and other forms of automated instruction, to erect a cybernetical theory of learning.

How shall we go about this? We shall have to call in the aid of the computer people. In the more sophisticated types of instructional programmes, students are led to take different routes or branches through the material to be learned according to their needs. But, at the moment, the only parameter on which this branching decision is made is the current response of the student to a current question. This is clearly inadequate. We would have the decision procedure take into account the student's past performance in the current subject and, indeed, in allied subjects. We should like to take into account his own assessment of his understanding. We would like the decision function or functions which decide at every instance where next to lead the student to be functions of more than this single parameter. If this is to be done, we shall have to use computers to do the data storage, data processing involved in the evaluation of these decision functions. But we shall be able to use computers in another, but very closely allied way, to build up this cybernetical theory of learning. It is perfectly possible to programme a computer to turn it into an intelligent machine. This has been done especially by people like Simon, Newall and Feigenbaum who with a set of programmes and a special computer language, called I.P.L.V. were able to build up an intelligent machine, as a set of computer programmes called Epam. Feigenbaum was thus able to show that there is no need to assume decay in an associative memory to account for forgetting, but that the kind of forgetting that happens in human beings can happen simply because the associative links between item and item are heavily overlaid by other items. Supposing then we were to build up

a provisional theory of learning, from the relics of psychological theories, our own information and the information fed back from teaching machines, we could programme a computer to compare this theory with the data from real life, and then itself to modify the programme so that we embark on an iterative process with each iterate in the process being a closer approximation to the kind of theory we want to develop. This is a problem which is not really difficult in principle. It's the same kind of problem that we encounter when we try to program a computer to work out the procedure by which a given set of numerical data are obtained from a given set of input data.

What other problems do arise? It seems to me that we shall want to be quite clear, first of all that we are not automating the whole teaching process. We are automating those parts of the teaching process at which human beings are not especially efficient. We thereby release the human teacher to become the human educator. In order to do this, we have not only to build up a satisfactory meta-language, but we have to subject many of the concepts we glibly bandy about now to very serious critical scrutiny. For example, if you look at the technical language or technical jargon we tend to use when talking about teaching machines and programmed learning, we discover that it is a remarkable hotch-potch. It is gleaned from various partial "psychological" theories or semi-languages. We haven't an adequate meta-language in which to talk about the things we are going to do.

All this, of course, is the work for specialists. What are the rest of us going to do in the meantime? I would suggest that our work will be more than cut out. First, it would be advisable to get down to the work of writing programmes instead of talking about writing programmes. There are the problems of the administration—problems that arise when we attempt to introduce programmed learning into educational and training establishments. One of the big factors to be taken into account here is, of course, the pace at which learning takes place. If we were ever to start with anything like a homogeneous group at the beginning of term and use programmed learning, we should certainly not have a homogeneous group at the end of term; in fact the concept of programmed learning and teaching machines cuts straight across the conventional idea of the class and the classroom. Indeed, I would go so far as to say that the teaching machine spells the end of the class and the classroom as we know it. Many of us will be sincerely thankful for that. Thirdly we have got the job of ourselves experimenting with, testing and validating our own programmes. Here we must emphasise that a teaching machine programme is not another form of text-book. The authority of a text-book derives largely from the author who has written it and the opinions of reviewers who may or may not have read it. But the authority of the teaching machine programme derives from the fact that it, like an aircraft, is a scientific product. It has been designed to do a job and it has been tested and modified and tested again until it is established that that job is done. If we are to do this successfully, we shall require the assistance of the statisticians among us. Many of the works on programmed learning that are available at the present time are, I would suggest, highly suspect from the point of view of statistical design. Many of them are concerned with only very short sequences, 10 frames or something like that. Many of them blithely ignore the existence of at least 210 of the 212 other variables involved which Fry has counted. In point of fact, what usually happens is this—that, appalled by the enormous number of variables that you have to deal with, you reduce your problem to apparently comfy two-variable problems, blithely forgetting or even not knowing, perhaps, whether or not the other variables interact. Then you test for this and you test for that and you come out with some startling paper which usually says very little; then you take another two-variable problem and another two-variable problem, and then you are optimistic enough to think that you can try to synthesize these results and get a solution to your multi-variable problem. And of course, you can't. We do need in all this the very serious advice and assistance of the statisticians among us. Probably one of the most useful teaching machine programmes that could be written would be a teaching machine programme in the main elements of the statistics of experimental design.

I would like now to come back again to certain more immediate aspects of the subject which tend to be disturbing. It really links up with this question of whether we are trying to simulate the "classical process of individual tuition". If you look at the dust-jacket of certain U.S. programmed text-books, you find this kind of thing said: "Computer programmers are at a premium, read this programmed text and you will command a salary of an enormous number of dollars a year", "English punctuation pays cash", etc. The advertisements that are creeping in as a result of its commercialisation are beginning to get extremely disturbing. All sorts of claims are being made for programmes which, quite frankly, they can never fulfil. So the question arises of the quality of the programmes we write and how we can tell whether the programmes we buy are of good standard or educationally sub-standard. For with the growth of commercialism, we can

be quite certain that a whole lot of programmes will be put on to the market which will be educationally sub-standard. How are we going to tackle this? It seems to me that we must insist, first of all, that each programme shall be prefaced by a number of things: First, there must be a detailed statement of the *criterion behaviour*, that is the explicitly stated objective of the programme. This we call the *criterion behaviour*, rather than the *terminal behaviour*, which is the performance of a student after he has worked through the programme. Secondly, it is absolutely essential that we should state quite specifically the entry conditions to the programme. This is not quite so important in the case of branching programmes which cater for a wider spread, but it is certainly necessary for the more linear types of programme. We should then preface each programme with an account—a detailed account, if necessary—of the procedure and the results of the tests which have been made on it to leave it in the form in which it is presented to the public. In this respect, may I draw your attention to a Temac programme recently produced—an *Encyclopaedia Britannica* programme—which though it doesn't go the whole way certainly is exemplary in these two respects. The programme on *Matrices and Mathematical Programming* has a statement of programming principles, a statement of the aims of the programme, a statement of the entry conditions of the programme and a summary of the tests and results achieved in those tests which caused the programme to be published in the form in which it is published. Not adequate in all ways, but nevertheless a very great step forward. Another Temac programme (and I am not advertising here), that on the *Algebra of Fields and Ordered Fields* is exemplary in the sense that for each section of the programme a teacher's guide is provided; this states categorically the criterion behaviour which it seeks to develop for each section of the programme. Now will this be sufficient? It seems to me that the time is coming when we ought to raise the question of the responsibility of professional organisations in this respect. I suggest for instance that where you have programmes which are in the province of a particular professional organisation, this organisation should help see to it that the programmes put out on the subjects with which it is concerned are not in fact sub-standard. If you have a computer society, for example, that society might surely look at and test programmes on computing. If you have a mathematical association, then it would be worth while for that association to concern itself with the quality and testing of mathematical programmes. In this way we should be able to eliminate ultimately, the sub-standard commercial programmes which are likely to be unloaded on us. These, then, are some of the problems with which we are confronted. May I make one final point? One importance of programmed learning, perhaps, is this: although we may never use teaching machines on a mass scale (although I think we shall), the very fact that we have invented them and are now forced to think about them means that we are also forced to think, maybe for the first time, about the fundamentals of what we are doing. And this is rather like dropping a pebble into a pond. The harder you think about the immediate circles, the more you are forced out to think of others. And so it may turn out that the teaching machine is a very stimulating and important device which has started us thinking once again about the things which we ought to have thought about for a long time, but, for various reasons, have not done so.