

# Some thoughts on Programming

By

PROFESSOR L. STOLUROW

I THINK we should look at the prose style used in programming as well as at the technique, because we are interested in trying to develop an efficient prose which will communicate and instruct. So it is important, I think, to realise that the kind of prose we are developing in programmed instruction is somewhat different from the prose that is typically developed in textbooks. The textbook really has two important factors, which act in a contrary way and have contrary goals—goals which are not instructional ones. One is that the book has assumed the general form and character of a repository, of an excellent reference; the textbook is a good place to find information when you know what it is you are looking for.

Also the prose book is designed to communicate to the *knowledgeable* person, whereas one thing in organising a strategy or in developing the sequence in a programme that one soon realises is that the organisation of the content has to be an organisation which clearly recognises the fact that the student is not an expert when he begins the programme, that he has very limited knowledge, and that we want to make him an expert by the time we get to the end. Books are typically organised from the point of view of the expert's perception of the material—and there is good reason for this. Text-books are not bought by the consumer, the student, but rather books are bought by college people who also teach. Education is probably the only industry in which the consumer does not actually determine the choice of the thing he will consume. When you write a book for a colleague, for an expert, you are obviously going to organise it so that he will see in it the kinds of things that he wants taught. When you programme, you are not bound in this same way; in fact the binding is quite different and somewhat more difficult, because what you are trying to do is to determine what the level of knowledge is of the individual entering that programme and then what you want to do with these various devices is to move him on up to the point of being an expert. One thing we try to do is to assess the level of information that the student has *before* we start giving him the programme; we identify the assumed entry behaviour, so to speak. Given

that he has not the required behaviour, then what we have to do is to give him some supplementary kind of instruction, maybe another programme, before he is put on this one.

Every time you sit down to write a programme, you have to keep in mind the two ends of your trip. You have to keep in mind the goals to which you want the student eventually to move and you have to make some assumptions about where he starts. Then you have to play a kind of game. You do this much as you would play a game of chess. What you do is you have this imaginary student go out in front of you and you ask yourself, from your knowledge of teaching, the kinds of things that might be most appropriately presented first; sometimes there are natural priorities, required priorities; many times there are choices, and here the teacher or programmer is required, of course, to make decisions on the basis of what he or she perceives as being the appropriate thing. Sometimes appropriateness is determined by just a limited view of the next few steps and you find after writing step 20 that you wish you had introduced something back in step 5 and you have to scrap the whole business and go back to 5 and start all over again because you didn't lay the ground-work sufficiently well for step 20.

This is the kind of cut-and-dried procedure which frequently goes on in writing programmes. Once you get something which seems to be sufficiently cohesive, sufficiently useful, then the next step is usually to bring in a student—and here's where you again get crushed, because you are quite amazed that the brilliant ideas that you really thought were going to get this notion across suddenly fall flat. The student just doesn't respond to them at all. But it's a very salutary experience and if you've got enough callouses and you are strong enough and prepared for this, after a while, you begin to appreciate its value and you begin to see wherein it can contribute to the overall effort. After all, we are developing now a prose form, the sole purpose of which is to teach and if this is the person who is being taught and he is having trouble, then you want to know it and you want to find out what to do about it. Usually

some kind of discussion takes place. You don't write a lot of frames. You don't wait to complete a whole programme, so to speak. You write a sufficient number to get a concept across and then you sit down with the student, let him go through it and then discuss it with him, find out where difficulties occurred, and look at the answers he wrote down. Constructed response is frequently a desirable thing for this purpose and the discussion frequently can be quite helpful.

Then you go back to the drawing board, revise and start anew, or you may save a certain set of frames. Frequently it is a good idea to prepare this material on cards. These don't even have to be typed often, they can be written out by hand. The idea is to put them on cards so that you are flexible, so that you don't have the difficulty of scrapping a lot of good items just because one part is bad; you can cut and splice, and students are usually quite tolerant of these things. The main idea, though, is to give yourself the kind of flexibility that you want. You can use a flash-card system with the right answer on the back and the student given just a pad of lined paper and an armchair to sit in and work; by these devices, you gradually build up to the point. Then you feel that you might be willing to tackle a larger group and you might want to feel that the investment in duplicating would be appropriate and desirable.

We, typically, then, in our pattern of doing things, get a larger group, maybe a class or so together and try it out, and again look at the individual responses, look at the time it takes and have some discussions with some of the students and then go back to the drawing board again; then it usually goes out to a larger and somewhat more extensive use.

This is some of the pattern of employment of these techniques. Frequently the students come up with suggestions that are quite helpful and useful in giving you ideas as to how to express a concept in a way which you hadn't thought about. I think their early experiences can be particularly useful when you are dealing with concepts that normally you don't have an opportunity to interact with students about. Conceptual sequences, in particular are important here.

I would like now to bring in some points about technique and get us beyond this procedural kind of problem I have just described to the questions again of what we might do and the kinds of things we might be concerned about when we do prepare these materials, these cards, in either the linear or branching form. So far as branching is concerned, there are some studies which suggest that, if it is used without selective

attention to the particular concept or the particular students for which it is being used, it may come out no more effective than a linear programme. In other words, what I am saying is that the studies that have compared branching and linear forms have not compared them when there has been some overriding consideration that one would be more or less appropriate than another for the material. When they are compared under these circumstances, we do not find that there is necessarily a consistent result with respect to their respective effectiveness. One study for example using material about the United Nations and another study using material about the number system, tested them by both bypass and linear type programmes. Usually these come out to be somewhat different in length and there are variables that are hard to control to make them identical. But, if anything, one gets equivalent achievement, usually in studies of this kind. Sometimes there are time differences, however.

We find typically the time differences tend to favour the branch programme, as you might expect: the bright student can flip through such a programme while the one who is having difficulty has to go into these branches and they take more time and, consequently, he takes a little longer. But, obviously, if they are effective remedial branches, when it comes to the end test, he should do relatively as well as the person who didn't need to know that material. So having them come out equal in the end is one way of indicating that the branch has done its job. One needs a little more sophisticated design, however, than the simple branch-linear comparison in order to really be able to assert that remedial sequences are in fact remedial in their function, and this is something which, unfortunately, has not been studied extensively. We are currently doing some studies on this problem, in other words getting at validity questions, questions of whether or not, remedial, so-called remedial branches, really are being used in that way. We find sometimes our guesses as to the need for remediation wind up again on the scrap heap because all the students or none use a branch. This frequently happens. Either it is a branch that no student goes into or it turns out to be a branch that every student goes into, whether he is a bright student in terms of aptitude score or a poor student. So judgment as to what constitutes a basis for branching is still an art form and still something that needs empirical check, close scrutiny.

In respect of review sequences, sometimes branches are introduced to review for the student the material he has been through. Here studies have shown that it is better to repeat missed items which the student did fail before he is allowed to

go on rather than at some later point. There isn't much data here, but those data we do have seem to make good sense and seem to be an appropriate finding, so even though meagre, I think they are worth mentioning.

With respect to sequence, we get into some interesting problems. Skinner, you know, added to Pressey's original contribution of the teaching machine, the notion of the sequencing of frames. Obviously the concept of shaping of behaviour would lead one to be concerned about the problems of how you lock things together in the way Mr. Goodman drew them on the board. Consequently the order in which you put your frames would be an important consideration. Up to that time the work done by Pressey and his students had not been with programmes for they had not developed programmes in that sense. Their "programmes" were a series of questions which queried the student in a not necessarily systematic manner and, based on his failure to answer them, the student was usually told to go back to an original reference source book. Once you introduce the question of sequence, the related issues arise, namely what kinds of rules do we follow. It is very commonly said "Oh, one uses a logical sequence, of course". Well this is a very tricky point because one will find that there is not one logic to any particular subject matter. I think that this becomes quite apparent when you look at books covering the same topic. Not all authors use the same sequence. In other words, for a given subject matter, it frequently happens that there are many sequences possible, all of which could be defended on the grounds of being logical. Consequently we cannot assert simply that a sequence is good simply because it is logical. One finds that alternative logics can be applied.

Naturally I am biased here. I would say that psycho-logics should be applied, namely that we should look at the problem psychologically to the extent that we can and that we should try to determine what it is that in the mind of the initiate we need to present in a particular sequence. This may not be the most logical sequence from our point of view. Mager, in California, has been doing some interesting things along these lines. He has, for example, sat himself down in a chair with an individual in a tutorial situation and he has said "Now I would like you to discuss with me some of the basic concepts in electricity". And he has said, "I know what you probably will say"—this will probably be a housewife or someone who is relatively unsophisticated with electricity—"I want you to tell me some of the things you know about electricity and then I want you to ask me some questions when you get to a point at which you feel that you are inadequate and I will try to answer them for you." So he starts

the person off and they get to a point and they ask a question and he answers it and so on, and he records this. Then what he has done is to build a sequence based upon the sequence of questions asked by the learner who is trying to discover for himself what this is all about. These sequences have very peculiar logic to them, if you can say they have a logic at all. Nevertheless when given to other students as a teaching sequence, they are effective. They seem to fit the requirements of other students and of other people equally unsophisticated in the area. This is by the way to illustrate a technique and to indicate that we are really just scratching the surface on the question of sequence.

We have been doing some things from a different point of view relating to this problem, many an approach which I think has some merit and, hopefully, will get us along in effecting a rapprochement between aptitude testing and learning. What we are doing is giving a battery of tests to our students prior to their taking a programme. These are aptitude tests and general ability tests. We have a distribution of scores on these. What we do is to divide them randomly into two groups. One group will get one sequence of steps, and the other group will get an alternative sequence. At the end we give them all the same achievement test. We also give them, sometimes, if we can get them together and reassemble them at some delayed point, a retention test as well. Then what we do is to correlate their scores on those various tests that we gave as our aptitude battery, with their subsequent achievement and retention scores. What we have found is that the achievement scores at the end for a group, the group means, are typically equivalent, indistinguishable, statistically drawn from the same population as far as we can tell. But what we find is that the pattern of significant correlations one gets with the aptitude scores is frequently and usually very different. This goes back to an analogy that Mr. Goodman used in his first presentation. You might think of this whole business of going through a programme much as you would a trip. The trip has a certain goal. The sequence of steps in one order by means of which you achieve your goal represent different aptitudes, apparently, from the aptitudes required when you take that same trip, so to speak, when you achieve that same set of goals, in terms of knowledge you acquire by following a different sequence of steps.

We have two or three studies in this connection one study with retarded children using a fractions programme, another using an earlier version of that same programme with normal fourth and fifth grade students, and we have some other data with high school, very gifted high school

students. All of these tend to show this same kind of distinction in the same way that sequence can make a difference with respect to the aptitude demands it makes on the learner. One other important thing which came out of these analyses, or seems to be coming out—I speak very tentatively because we need to do very many more studies and some further analyses—is that we can by the very way in which we sequence the material apparently trade aptitudes or abilities; in other words we find that with one sequence, mental age, or the kind of thing that goes in as a basic factor in determining IQ, can be reduced to an essentially zero correlation with final achievement by one sequence and can be significantly related to final achievement with another sequence. Now while those shift in that direction, we find an opposite kind of shift with respect to specific aptitudes in these mathematics or arithmetic programmes. We find, for example, that specific mathematics aptitude, arithmetic comprehension and so on will be significant correlants of final test performance when the MA is not significantly correlated, and *vice versa*. This tends to be the way in which the balance swings. You trade off the general ability significance for more specific prediction on the basis of particular aptitudes.

If this tends to hold up then, of course, it moves us on to the direction in which we hope to be going, starting next September, namely to investigate some problems in this whole area that we haven't even touched upon, which I think are a part of it, and which become essentially a part when one has a computer-based system available. This is the pre-tutorial decision-making aspect of the whole business.

We have been talking about the tutorial instruction up to this time and most people in our country have only been concerned with this too, but we feel that there is a definite need for formalising the decisions one makes which relate to the profiles of ability of the learner to the particular strategy of instruction one uses when he goes about teaching.

The idea is that we have made very excellent aptitude tests. We have some very good instruments that give us interesting profiles that have high validity and reliability. But I defy you to point to this many instances in which these ability tests have actually been used to specify the differential way in which instruction should be accomplished. This I feel is a serious problem, and this is one of the problems we hope that the methods we are talking about at this Symposium will move us on to, as another one of the frontiers. There is obviously a relationship between aptitude and the kinds of approaches one would take to the instruction of a student; one should examine

these pre-tutorial decisions just as one examines the actual strategy that one uses once one has decided on an analysis of the goals and objectives and made some assumptions about initial knowledge.

I won't go into the formal analysis, but we have made an attempt to develop this, by making this kind of decision automatically on a computer into which we feed a profile set of scores. On the basis of its programme, the computer will make decisions about whether to use strategy A, B or C. We shall make lots of mistakes; we shall pick wrong strategies out; we shall have to build this information up by a bootstrap so to speak. We are trying to do this now, prior to the actual use of the computer, using human individuals to simulate the computer as it were, and making these decisions on a preliminary basis so we shall know how to begin.

The other end of the line that I simply want to mention before stopping is that we are examining the question of the kind of decision one makes to improve one's strategy while instruction is going on. Most of the procedures that we are able to accomplish now with the devices you see here, no matter how elaborate they may appear to you, are, nevertheless, quite restrictive when one gets into a very general concept of what one might like to do in the best of all possible worlds. For example, it may be possible to give a student a series of problems in which he is applying a new concept. You give the student a variety of different problems bringing in different sets of conditions, and the computer could be observing not just each of these as individual responses, but it could look at the pattern or set of them, and it could have a programme in it which would say that if any three out of four of these are wrong solutions, then you take the individual back to point  $\times$ , and, if all of them are right, you skip him ahead to  $\times$  plus 20 and so on. In other words, one would make much more elaborate use of the branching idea and base it upon information we are not now able to use in making branching decisions. We all know that decisions of this kind are extremely unreliable if the information base is meagre. Consequently, one of the directions in which I think we need to move is to fortify or build up the basis for decisions about branching. The two directions in which we would like to move most immediately are to utilise aptitude data in consort with data about the immediate history of the individual going through the programme. We call this procedure dynamic branching or dynamic decision-making.

Well, this is the current direction—we are now off into the future. I might, in the brief few moments available to us, sum up with a couple

of points. I think that our needs immediately—and I hesitate to say this to be prescriptive for use since I have been here for a matter of less than 100 hours I believe—but it seems to me that one of the problems to hit hard with programmed instruction is the problem of teacher-training. I think what you need to do is to focus on your teachers first and your pupils second. In other words, I think you want to try to involve as many teachers in the whole operation as you can, to see how many of them might eventually contribute in the sense of developing material themselves, because, obviously, one wants to build up capital, a resource which would multiply and proliferate. If one applies the method directly to the lowest echelon, I think its impact might be somewhat less than it would be if it were applied, let us say, to the middle group, the group which might, as a result of such experience, improve their teaching. Incidentally we find some of our most able teachers who get new insights into their subject simply by trying to programme it. I think some of these teachers would become useful programmers; and one thing you need, of course, in order to get the show on the road is to have programmes. All of the instrumentation in the world of the kind we see here, no matter how good and how elaborate and how reliable, is not going to do the job unless you have the programmes to feed it. I think that we need to lay the groundwork by developing conferences of this kind in which possibly some of you will go back to the groups with whom you work and describe the general approach and tell the people who will eventually be involved in it what it is about to acquire information.

It think there is a preparatory stage that's needed in order to get wide acceptance and use. I think there are many cautions that have to be considered. One of them is this cultural caution that I mentioned earlier. I think one has to import. Obviously one can't start from scratch. But one has to be careful. One has to examine programmes, try them with small groups, get reactions to them and get some feeling on a very personal basis as to how they are working and what kinds of difficulties are developing, going through some of the steps I described a moment ago. I think that, in doing this, you want to develop and try to standardise, if you can, your interrogation procedures, so that you get useful information fed back to you. Then you want to move on to a little larger group usage, collecting data and examining responses. I think you want to be careful of aversive stimuli creeping into programmes. There is a great tendency for teachers to tell students when they are wrong that they are dumb or that they picked a very stupid answer, etc. One of the impacts which that kind of aversive

feedback has is to eventually discourage the student from learning. It is probable that very good teachers frequently are autocratic and good students will follow along. They get meted out some punishments, but not too many. The poor student who can't keep up the pace gets many. Consequently, he not only doesn't learn that subject but he also learns not to like learning, and this is the tricky part. Aversive stimulation tends to make the individual avoid the learning situation, and what frequent reinforcements tend to do is the contrary, namely to get him to like the learning situation, and sometimes, if they have no other impact on the learner than that, they should be employed so that you get the mechanism going so to speak.

So I think one caution then is to avoid aversive stimulation. Here cultural knowledge is extremely important. We have toyed, for example, with the idea of telling an Indian in India that his response was "fat" because we know, from analysis of the language, that obesity in our culture would be considered to be important, high valued, even the rich can't afford to be fat, you see. In India things are so meagre. I am using this as an illustration to indicate that one wants to consider the kind of language that one uses in providing knowledge of results to the learner in terms of its cultural impact, and I mean sub-cultural as well as general culture. There are all sorts of sub-cultural groups within any culture and they tend to value things differently from, let's say, the teachers who are in that culture. Consequently, the kinds of rewards which the teacher thinks he is meting out to the students aren't necessarily perceived by the learner as being the kinds of rewards he would like. So we need to be cautious in this sense in providing information feed back and seeing that it is suited to the individual learner we are dealing with. It may turn out that important programmes are modified most in the feed back information that's given the learner more than in any other aspect.

I would like to wind up on the note of assessment and a caution and an emphasis on the importance of tests. We have talked about programmes and we have said how good they are and what they can do if properly developed like any instrument or concept. If it is properly developed, it can have lots of pay-off. But we have not said anything about how we determine that pay-off. Of course the measurement of performance at the end is a very critical aspect to the whole business. One looks at a programme and is sort of dazzled by it because of its newness and so on. But I hope that you will not be misled into thinking that that is all. One needs also to look at the examinations that are given following the use of the programme really to get a notion as to what it is that that programme

is teaching, the real pay-off, then, in the sense of the data I have told you about to-day and in the sense of the manuals that will accompany programmes of the future and tell you how effective they are. The important thing is that information plus a knowledge of the test.

Testing the assessment of programmed achievement is a really tricky business and more here than meets the eye. The basic dichotomy that one wants to think of in this connection is the dichotomy between the normative testing, which is the predominant form that we have with us to-day, where tests scale individuals relative to other individuals, and this is what you find most tests doing, and programmed instruction. What we want in programmed instruction and, I think, education generally, are not tests which scale one individual with respect to another, but tests which inventory, tests which determine the

particular components of knowledge or skill that are lacking. We want to do first a task analysis, find out what it is that the individual needs to know and then we need to design a test to cover each of these. We have to set up a kind of matrix, you might say, of information concepts, let's say, and applications, and you might set this matrix up in different ways and then you can visualise it as being a set of cells. What we need to know is whether the individual knows what's in cell 1, cell 2, 3, 4 and 5. If what he doesn't know can be identified in this way, then portions of the programme can be related to that lack of knowledge and he can be put on the road, so to speak, in a minimal amount of time. If, on the other hand, we have tests that simply scale individuals relative to one another we really do not know where specific deficiencies lie. There is a whole new area of development here.

---

## WILEY BOOKS

### A GUIDE TO PROGRAMMED INSTRUCTION

**Principles and Practical Techniques to Programming Instructional Units.** *By* JEROME P. LYSAUGHT, University of Rochester, and CLARENCE M. WILLIAMS, University of Rochester. By offering step-by-step guidance, this book presents a practical process for developing a programmed instructional unit. It shows how to select appropriate units of study for programming, how to design programmed materials, how to utilize them, and how to introduce them into teaching programmes. The simple statistics offered here make the book a self-contained guide to programme construction and utilization. The illustrations have been drawn from the work of actual programmes, and they clarify each step of the programming process. 1963. 180 pages. R3.35.

CHECK ALSO:

Coulson (Ed.): *Programmed Learning and Computer Based Instruction*. 1962. 291 pages R5.75.  
Margulies-Eigen (Eds.): *Applied Programmed Instruction*. 1962. 387 pages. R5.95.  
Galanter: *Automatic Teaching: The State of the Art*. 1959. 198 pages. R3.20.

*send now for your on-approval copies*

JOHN WILEY & SONS Inc. 605 Third Avenue, New York 16, N.Y.

C.N.A. - AGENCY DIVISION 14 Nelmay Court, 6 Ameshoff Street, Johannesburg