In Don D. Bushnell and Dwight W. Allen (Eds.), The Computer in American Education. New York: Wiley, 1967.

chapter 2

On Using Computers to Individualize Instruction

Patrick Suppes

The theme of individualizing instruction is a very old one in education, and the supporting psychological documentation of significant individual differences in initial abilities, in rates of learning, and even in general approaches to learning, is by now overwhelming. Yet it is fair to say that for simple reasons of economics we have not been able to individualize instruction to any very deep extent at the elementary, secondary, or college level. Costs for providing this sort of instruction are simply too great for a society like ours, committed as it is to universal education at all but the highest levels. Administrators and teachers continue to struggle with the problems of meeting every child's educational needs. In the elementary school serious attempts are made at grouping within the classroom for important subjects like reading, and recently the ungraded elementary school has begun to attract considerable attention. Without question the main intellectual justification of the ungraded classroom is the opportunity it provides for a greater degree of individualized instruction. To some extent the introduction of team teaching has had the same sort of motivation, although in this case another aim has been to provide better qualified teachers by appropriate pairing in skills and training.

Efforts like the development of ungraded elementary schools undoubtedly will continue, and will be important in the future as we make further attempts to tailor education to the individual child. Other solutions are clearly needed, however, if we are to approach anything like a tutorial mode of individualized instruction. It now seems that the burgeoning technology of computers may offer a real avenue of approach to providing the kind of instruction we want.

The application of computers in school instructional programs has barely begun and is not yet well understood, but some examples will be sketched in sufficient detail to give the reader a feeling for the real prospects and concrete problems of this sort of instruction. Since my own efforts have been devoted mainly to the elementary-school mathematics curriculum, my examples will be drawn primarily from this part of the curriculum. One or two other specific possibilities, such as foreign-language teaching at the high-school and college levels, will be mentioned.

COMPUTER SYSTEMS

THE "LOOK" OF COMPUTER-CONTROLLED TEACHING DEVICES

There are three levels at which computers may be used to individualize instruction, but perhaps it will be helpful to describe first the kinds of devices or terminals that may be attached to the computer for students to use. Of primary importance is a visual-display device on which the student may look at messages brought up from computer memory. A device that is now becoming familiar is the cathode-ray tube (CRT), which looks very much like a small television screen, and on which alphanumeric character messages may be generated directly by the computer. Attached to the CRT is a standard keyboard that the student may use to respond to the instructional program. In some environments the student will also have a light pen that he may use to point to correct responses on the screen of the CRT. A device such as a light pen is particularly important for young students in the elementary school. Communication with students through appropriate audio devices is desirable at all levels, though again it is especially needed with younger children. With the more developed programs, prerecorded messages are available fairly rapidly at appropriate points. Additional visual-display devices are also important; for example, displays of photographs or more complex line drawings than are easily presented on a CRT are useful at almost all levels of instruction. Since the CRT equipment will make no appreciable noise while in operation, the devices described here can be placed in an ordinary classroom if the student uses earphones. If the CRT is replaced by a typewriter or teletype, then some sort of acoustical isolation is necessary for ordinary classroom operation. A variety of configurations seems practical, ranging from a single CRT in a classroom shared by many students during the day, to a classroom containing a large number of CRT's. In the latter configuration each child would spend at least an hour a day on the CRT, perhaps working in thirty-minute sessions, once in the morning and once in the afternoon; but it should be emphasized that he would spend most of the time in his regular classroom. Again I have in mind the elementary school; the appropriate change in environment required for college applications should be obvious.

There is already a good deal of evidence that students at all age levels gradually come to feel as much at home with the sort of terminal equipment described here as they do with an ordinary television set. Of course, it is understood that young children in the elementary school are not expected to input long messages on the keyboard. For most exercises a few digits corresponding to numerical answers, or a single word, would be the most required. Longer responses are appropriate for older students; we shall return later to the rather complex problems of evaluating constructed responses of any substantial length.

In current work at Stanford University on computer-assisted instruction, there are three levels of interaction between student and computer program. A brief description of each level will be given, although it must be emphasized that the third and deepest level is still mostly beyond us from a technical standpoint. For now, it represents an ideal to be attained in years to come. Each of these levels will be referred to as a *system*, in conformity with rather widespread computer usage. The point is that in each case a computer system is built up to perform the required level of instruction.

DRILL-AND-PRACTICE SYSTEMS

This level of interaction is merely supplementary to the regular curriculum taught by a teacher. In the case of elementary-school mathematics there is abundant evidence from both pedagogical and psychological studies that students need a great deal of practice in the algorithmic skills of arithmetic before a reasonable level of mastery is obtained. They need corresponding practice in the standard applications of arithmetic and, more generally, in developing what is sometimes called a good "number sense." The point of a computer system at this level is to provide a simple, straightforward, and *individualized* approach. It is intended to relieve the teacher of a considerable burden and at the same time take a substantial step toward providing practice work at a level appropriate to each student.

For the minimal drill-and-practice system we can forego the audio component of the total system described earlier, as well as the light pen. We can use a CRT with keyboard, or a typewriter under computer control with keyboard. The CRT makes it possible to display multistep algorithms like that of long division in a more direct fashion than the typewriter allows. Even with the CRT and keyboard, we might make the response of traditional education and give elementary-school students a standard set of exercises that will be the same for all, with the familiar ring of drill and more drill. But the most important difference from traditional methods is that we are not committed to giving each child the same problems, as we would be if textbooks or materials prepared at the school were used. It is a straightforward matter of computer programming to offer exercises of various degrees of difficulty, and to select each student's level according to his past performance. In a program we are now running in the upper-elementary grades at Stanford, there are five levels of difficulty at each grade level and on each concept in elementary-school mathematics. For instance, five different exercises illustrating the fact that subtraction is not commutative might be given to different students on the basis of their past performance, at the same point in the school year. Typical of such exercises, which also provide additional practice in the algorithm of subtraction, are the following:

A. $4 - 2 = 2 -$ _____.
B. $16 - 12 = 12 -$ _____.
C. $34 - 25 = 25 -$ _____.
D. $63 - 45 = 45 -$ _____.
E. $123 - 75 = 75 -$ _____.

A moment's inspection shows that the five exercises all exemplify the same general concept and yet vary considerably in levels of difficulty, particularly in terms of computational skill.

To take another example, a typical three-day block of problems on the addition of fractions would vary in the following way. Students at level A get problems that involve only fractions having the same denominators. Levels B and C, on the first two days, also have only problems in which the denominators are the same. On the third day the fractions have denominators that differ by a factor of 2. At level D the problems have denominators that differ by a factor of 2 on the first day, and at level E the denominators differ by a factor of 3, 4, 5, or 6.

At the present time we are moving the students up and down the levels of difficulty on the basis of the previous day's performance. If more than 80 per cent of the exercises are correct, the student moves up one level, unless he is already at the top level. If less than 60 per cent of the exercises are correct, the student moves down a level, unless he is already at the bottom. If his percentage of correct answers falls between 60 per cent and 80 per cent he stays at the same level. It should be emphasized that the selection of exactly five levels and of the percentages 60 and 80 has no firm theoretical basis but is based on practical-pedagogical judgments. As systematic data are accumulated, we expect to modify our choices in the light of experience. The important point for the present is that we are operating a highly individualized program of instruction, which even at the level of drill and practice depends on the availability of computer-based terminals. In principle the work we are doing could be done, of course, without a computer, but only with a very substantial addition to the teaching staff. It would hardly be possible in practice. Obviously this is true of almost everything that we ask of computers, not only in education but in industry and government as well. The operations carried out by the computer could always in principle be carried out by hand, given enough time and personnel. In the drill-and-practice exercises we are discussing, the computer can be used to analyze and collect data in a fashion that would be extremely difficult for a teacher. Above all the computer can systematically make an item analysis as well as present a daily written record on each student. The particular significance of the item analysis is that it can be used to refine the selection of items for the different levels of difficulty, and can give the teacher information on the kinds of concepts that are most difficult for her students.

There are also deeper levels at which drill systems can operate. I have described the preparation of five levels of exercises, which are prepared in advance of any use by the student. It is also possible to write programs with a more complex structure in which the individual items presented are contingent upon previous student performance. That is, items in each category are selected on the basis of a more detailed use of a student's past history than the one just described. For example, a student may be given more problems that emphasize the types that he has found most difficult over the past week or two.

TUTORIAL SYSTEMS

In systems of this type, in contrast to the drill-and-practice systems, the aim is to take over the main responsibility for developing skill in the use of a given concept. In the teaching of elementary Russian at the college level, for instance, it could be the responsibility of a tutorial system to offer a complete body of curriculum material on phoneme discrimination during the initial hours of the course. It is also easy to select an example from the initial lessons in arithmetic in the first grade. It is a familiar experience of many teachers that children entering the first grade cannot properly use the words "top," "bottom," "left," "right," and so forth. In order to give instructions it is highly desirable that the meaning of these words be clear to the children and that they be able to respond in unequivocal fashion to instructions using them. Here is the sequence of concepts we have used at Stanford in the first section of first-grade arithmetic:

1. Using the light pen—the child uses the light pen to point to the picture of a familiar object shown on the CRT.
2. Using a box to respond—the child touches the light pen to a small square box shown on the screen next to a figure.
3. Auditory introduction of the words "first" and "last"—the computer speaks to the child.
4. Auditory introduction of "top" and "bottom"—the kind of instruction given to familiarize the child with the use of these words is of the following sort: "Put your light pen on the toy truck shown at the *top* of the display."
5. Introduction of the word "middle" in the sense of vertical position.
6. Introduction of "left" and "right" in selecting one of two things.
7. Introduction of "left" and "right" in selecting among three things.
8. Introduction of "middle" in a horizontal sense.
9. Introduction of "left" and "right" in selecting among a row of things shown.
10. Introduction of "between" in selecting among several things—for example, "Pick the chair that is shown *between* two tables."

Space does not permit a description of the details of the tutorial program used to introduce these various concepts; but this concrete example may demonstrate that by using the auditory capacity as well as the light pen attached to the CRT, it is possible to approximate the interaction a tutor would have with a student. Above all it is possible to analyze each child's comprehension in greater depth and detail than is usually possible for a teacher of thirty students in a first-grade classroom.

It should also be apparent that in the tutorial system we may individualize instruction for the entering first-grade child. The bright, organized child who has been going to kindergarten and nursery school for three years before entering the first grade, and who has a large speaking vocabulary, could easily go through the concepts I have listed in a thirty-minute session. The culturally deprived child who did not attend kindergarten may take as many as four or five sessions to get through these concepts. These slower children may also be handicapped by an inadequate vocabulary in standard middle-class English, since their speaking experience may have been mainly with a dialect. These children often will require several different approaches before the concepts are thoroughly mastered. The important point is that in the tutorial program every effort is made to avoid an initial experience of failure for the culturally deprived child. The program also has enough flexibility to avoid boring the bright child with endlessly repetitive exercises that he fully understands. The child progresses through each concept in the sequence listed as he meets a criterion of performance for each. The child who makes no errors—that is, who fully understands the concepts to begin with—will meet each criterion very rapidly and move through the entire sequence in a single session.

In view of the kind of examples presented, the reader may feel that the instructional mode in a tutorial system will not permit freely constructed responses on the part of the student. Without entering into details here, a brief discussion of our program in mathematical logic will show that such responses are possible. This is one of our best-developed programs. The student is permitted to make any valid inference, and the main function of the tutorial program is to assess the validity of the inference he makes. The relative wisdom of the step taken by the student is not indicated until he has at least made a serious effort to find the correct proof of a theorem. As is usual in proofs of mathematical theorems, different students will find different proofs, and the computer program will accept any proof that is valid. When students do not succeed in finding a proof, the program gives them hints.

Concerning the future prospects for tutorial systems, it should be evident from the sorts of examples stated that the skill subjects such as reading, mathematics, and elementary foreign languages can be handled most easily and are best understood as taught in this environment. Tutorial systems can be used to carry the main burden of teaching skill subjects. Widespread application would lead to a radical revision of the organization of teaching, because a rather large part of all instruction is at the elementary-skill level. For this reason it is common to ask what will be the teacher's task if these elementary skills are taught in a tutorial fashion by computer-based terminals. The most important point to be made in this connection is that no tutorial program for computers in the near future will be adequate to handle every type of problem that arises in students' learning. It will be the teacher's responsibility to move to the much more challenging and important task of troubleshooting, of helping those children who are not proceeding successfully through the tutorial program and who need some sort of special attention. At Stanford University the tutorial programs have what is referred to as a teacher-call. When a student has run through all the branches of the concept, and has not yet met a satisfactory performance criterion, there is a teacher-call at a proctor station, and a teacher comes to give individualized instruction as extensively as needed.

DIALOGUE SYSTEMS

What we envisage as dialogue systems are computer programs and appropriate terminal equipment able to conduct a genuine dialogue between the student and the program. Such dialogue systems exist now only as elementary prototypes, and successful developments in any depth demand the solution of some relatively difficult technical problems.

Two central problems may be described by beginning at the college level and moving down to the elementary-school level. Suppose that in a program on American history, the student types in the question "Why did Booth kill Lincoln?" or a more complicated question such as "What was the role of the railroad in the economic development of the Mississippi Basin in the nineteenth century?" It is a very difficult problem to write programs that will recognize and provide answers to freely constructed questions of such generality and complexity. The situation is by no means hopeless, however. In curriculum areas that have been taught for a considerable time and that have a reasonably sharp focus of subject matter, it is possible to provide a fairly thorough analysis of the types of questions that will be asked. In these subject areas we can make considerable progress toward the recognition of the question by the computer program. The central intellectual problem at the moment is not that of writing information to give an answer, that is, of having in storage information that will give an answer to any question. Rather it is to recognize from the standpoint of the program precisely what question has been asked.

The second sort of problem arises in working with elementary school children for whom it is essential that we be able to recognize their spoken language. It is certainly not reasonable to expect a child in the first year of schooling to be able to input a question on a typewriter. He can ask or answer a question in a fairly complex way, however, if his speech can be recognized by a computer program. The problem of speech recognition simply adds another dimension to the problem of recognition of sentence meaning. There is reason to hope that within the next five or six years much progress will be made in the area of speech recognition. If the difficulties can be overcome, then in many respects the problem of the elementary school is easier than that of higher grades, for the types of questions and answers that occur are considerably simpler than at the secondary or college level.

PROBLEMS

Whether the computer system we are using to individualize instruction is a fairly thin drill-and-practice system or a very rich dialogue system, the kinds of problems that arise in designing curriculum materials and in organizing and using the system are very similar. In this brief survey emphasis is placed on behavioral rather than technological problems. Until we know how to settle the kinds of questions I shall raise, it will be unclear precisely what technological requirements must be met. To some extent the listing of behavioral problems is grouped around the classical concepts of stimulus, response, and reinforcement, but the classical formulation of these concepts certainly is not adequate to the problems we encounter in individualizing instruction.

STIMULUS SEQUENCE

Perhaps the first problem we encounter is the gross one of how we should think about organizing the curriculum materials. In what order should the ideas in elementary mathematics be presented to students? What mix of phonics and look-and-say is appropriate for the beginning stages of reading? Should phoneme discrimination be taught before word recognition or phoneme production in foreign languages? These are the kinds of difficult and perplexing questions that arise, even in the very early stages of preparation to teach mathematics, reading, or foreign languages. What is to be emphasized is that we are very far from a detailed scientific answer to any of these questions. Individualized instruction in a computer environment must proceed for some time on the same basis of practical judgment and pedagogical intuition that we use in arranging curriculum materials for ordinary courses in ordinary classroom settings. We have the hope of developing a more scientific and therefore a deeper understanding of these matters, but we are still far from having that understanding. The magnitude of the problem of stimulus or curriculum sequencing is difficult to overestimate. It is easy, for example, to lay out experiments that would need all the children in the world as subjects, in order for us to decide how to teach mathematics in the first grade. This problem of sequencing, magnified once again by relative priorities for different subject matters, leads directly into a combinatorial jungle that even large school populations and the power of the computer are not adequate to touch. It is a straightforward matter to show that the number of possible sequences of concepts and subject matter just in the elementary-school curriculum is in excess of 10^{100}, which is larger than even generous estimates of the number of elementary particles in the universe.

The only sensible hope for emerging from this combinatorial jungle would seem to lie in the development of an adequate body of fundamental theory about the learning and retention capacities of students. The history of science in many of its domains (particularly astronomy) is testimony to the importance of having large bodies of dependable data for theories to describe and explain. It is to be hoped that as systematic bodies of data become available from computer systems of instruction they will have a measurable impact on the development of learning theory. At the present time extrapolations from the kinds of experimental studies done to support fundamental learning theory and the central problems of subject-matter learning are tenuous, and existing theory inadequate for the complex problems of learning that arise in the context of the curriculum.

RESPONSE MODES

We do not yet know how critical it is to have various response modes available for instructional programs. The problem of interpreting complex constructed responses has already been mentioned; how essential such responses are to the learning of most elementary subjects is not fully known. A problem at least as difficult is to organize computer programs in such a fashion that the responses of the student are used in an insightful and informative way, both in telling him things and in deciding what he should be exposed to next. One of the most difficult tasks is to know how to make use of unexpected responses as a good tutor would. For the immediate future, perhaps the best use that can be made of unanticipated responses will be to record them for off-line analysis by those who must revise and improve the curriculum materials.

Again the problems we face are not easily overestimated, if we look ahead to the full task we would like the computer system to undertake. Beyond question one of the most tedious and unrewarding aspects of teaching English, either at the secondary or college-freshman level, is the thankless task of marking student papers. It is hard to think of a teacher task that one would rather turn over to a computer, particularly for routine mechanical errors. Unfortunately, apart from a check of spelling, it is still extremely difficult, and from a practical standpoint impossible, to give such tasks to any computer system now available.

REINFORCEMENT SCHEDULES

The possible types of reinforcements are a source of many open problems in designing curriculum for computer systems. A wide variety of alternatives are available to us. A difficulty is the lack of adequate theory to inform us on the best procedure to follow. There is conflicting evidence as to whether or not students should be corrected immediately each time they make an error. It is not clear to what extent students should be forced to seek the right answer, and indeed whether this search should take place more in what has come to be called a discovery or inductive mode, as opposed to more classical modes of instruction which consist of giving a rule followed by examples and then exercises or problems that exemplify the rule. A particularly troublesome issue that has come to the fore in recent research is the question whether different kinds of reinforcements and different sorts of reinforcement schedules should be given to children with different basic personalities. So far as we know, no large-scale curriculum efforts that build in variables of this sort have yet been made in this country. It would be very difficult to think about how to write two different elementary-school mathematics curricula with this kind of variable in mind—although some of my colleagues are more sanguine about it.

COGNITIVE STYLES

Closely related to the last point about personality differences is the body of evidence bearing on whether children have fundamentally different cognitive styles—for example, are either impulsive or reflective in their approach to problems. Perhaps the primary difficulty with the research on cognitive styles, as it relates to the construction of curriculum for computer systems, is that the research is primarily at an empirical level. The reach of theory is as yet very short, and it is not at all clear how the empirical demonstration of different cognitive styles can help us to design highly individualized curriculum materials adapted to these different styles. An even more fundamental question of educational philosophy asks how much the society wants to accentuate these differences in style by catering to them with individualized techniques of teaching. I do not mean to suggest that a curriculum design using information about these differing styles would necessarily cater to them, but this is one direction in which it is easy for curriculum to move, once such problems are brought to the surface. One may be reminded, in this connection, of the kinds of curricula that have been given children with special talents in our society for many years—for example, children with exceptional musical talent.

PAST PERFORMANCE

In the near future the use of computer systems to individualize instruction will surely encounter a more pressing problem than those of personality differences or cognitive styles: the problem of making effective use of information about the recent performance of the student. In standard classroom teaching it is impossible to use in any sensitive way the different achievement records of the students in the class. Partly because of the requirements of group teaching, we have very little experience in such matters. A gifted tutor will remember and use many facts about the past performance of his pupil, but scientific studies of how this should be done are as yet in their infancy. Practical decisions about the amount of review work needed, the time needed for the introduction of new concepts, and so forth, which vary widely from student to student, must ultimately be much influenced by the student's past performance.

DATA LOGISTICS

When students are set to work on an individualized basis, the problem of keeping records of work accomplished is enormously complex, particularly when it is intended that those records will be used to make decisions about the next stage of instruction. We have found at Stanford, in planning for the processing of 5000 students per day through a large computer system, that one of the most difficult decisions is that of selecting the small total amount of information it is possible to keep permanently. It is very easy to give teachers, administrators, and curriculum writers such an overwhelming amount of information that there is little chance of their absorbing and using it effectively. In terms of the variety of summary statistics one may easily obtain in a computer system from a day's work with students, it is not at all difficult to have the data output run to a thousand pages when 5000 students are being processed. The problem is to reduce the number of pages of data output from 1000 to something more like 25 or 30. The most essential problem in making this reduction is that we do not yet have well-defined fundamental theoretical ideas that can provide guidelines for making the reduction. Present decisions about what reductions to make are based primarily on general pedagogical intuitions and the scientific traditions of quantitative data analysis in experimental psychology. Neither of these guidelines is very effective.

It is easy enough to talk about individualizing instruction in the schools. This brief list of problems should convey how complex is the task of actually implementing such instruction by use of computer systems. Above all it is important to realize that the problems are not simply the technological ones of having computers available and adequate computer programs for preparing and presenting curriculum materials. Any deep and creative use of computer systems for individualizing instruction faces major theoretical questions still unanswered. From a scientific standpoint the search for the answers to these questions is the most exciting aspect of the whole venture.

CONCLUSION

Computers do not offer a panacea for the many problems of instruction we face in our schools. As with any new technology with a potentiality for solving old problems, the initial development and implementation create nearly as many problems as are solved. The road that lies ahead is tortuous and bumpy, but the signs point upward toward the heights of educational excellence. Many of those most concerned with the preservation of humanistic or individualistic traditions view the increasing importance of computer technology in our society as a threat to those traditions. If such critics understood better the scientific world in which we live, they would realize that this technology offers a possibility of individual fulfillment in education hardly conceivable fifty years ago. It is not too much to claim that for the first time since public education for everyone became a major goal of our society, individualized instruction at a genuinely deep level is now a feasible goal.