Teaching Children to Read Using a Computer

RICHARD C. ATKINSON *Stanford University* [1]

For the past 10 years I have existed in two quite separate worlds. One world is that of an experimental psychologist working in the isolation of a university laboratory on problems of memory and cognition; the other is that of an applied researcher attempting to computerize the instructional process. It would seem that there should be a fair amount of commerce between these two worlds, but I am disappointed to find that very few of my colleagues in memory and learning are aware of the work on computerized instruction, and this situation is equally true for my friends in education. Therefore, this article gives me an ideal opportunity to propagandize a bit for the potential that each of these fields has for the other.

This article is primarily descriptive, focusing on work that we have been doing at Stanford in teaching reading to first-, second-, and third-grade children. Let me emphasize, however, that there is a clear link between this work and basic research on memory and cognition. I like to refer to this link as a "theory of instruction." By that phrase I do not mean a highly formalized theory, but rather a loose collection of theoretical and empirical facts that can be used in conjunction with educational methods to design optimal procedures for instruction. Having written about the ingredients for a theory of instruction, I will not spend time on that topic here.[2] Simply stated, there are examples in which psychology provides powerful tools for devising optimal procedures, particularly when instruction can be brought under computer control. I will refer to several of these examples, but will discuss only two in detail.

[1] This article was an invited address presented at the meeting of the Southwestern Psychological Association, Dallas, Texas, April 1973. The preparation of the article was supported in part by National Science Foundation Grant NSF-GJ-443X3 and in part by Office of Naval Research Contract N00014-67-A-0012-0054.

Requests for reprints should be sent to R. C. Atkinson, Department of Psychology, Stanford University, Stanford, California 94305.

[2] See Atkinson (1972), Atkinson and Paulson (1972), and Groen and Atkinson (1966).

Computer-Assisted Instruction and the Reading Curriculum

Our first efforts to teach reading under computer control were aimed at a total curriculum that would be virtually independent of the classroom teacher.[3] These early efforts proved reasonably successful, but it soon became apparent that the cost of such a program would be prohibitive if applied on a large-scale basis. Further, it was shown that some aspects of instruction could be done very effectively using a computer, but that there were other tasks for which the computer did not have any advantages and possibly had some disadvantages over classroom teaching. Thus, during the last four or five years, our orientation has changed and the goal now is to develop low-cost computer-assisted instruction (CAI) that supplements classroom teaching and concentrates on those tasks in which individualization is critically important. A student terminal in the current program consists only of a Model-33 teletypewriter with an audio headset (see Figure 1). There is no graphic or photographic capability at the student terminal as there was in our first system, and the character set of the teletypewriter includes only uppercase letters. On the other hand, the audio system is extremely flexible and provides virtually instantaneous access to any one of 6,000 recorded words or messages.

The central computer which controls the CAI system is housed at Stanford University. Telephone lines link the computer to student terminals located in schools near the University and as far away as Florida, Oklahoma, Texas, and Washington, D.C. First-, second-, and third-grade students receive CAI reading instructions for anywhere from 15 to 30 minutes per day. Instruction begins with the student typing R for reading, an identification number, and his first name. The program responds with the student's last name and

[3] For a review of this work, see Atkinson (1968a, 1968b, 1969).

Figure 1. Student running on the CAI reading program. (The terminal consists of a Model-33 teletypewriter and earphones with an audio amplifier. The program operates on a PDP-10 computer located at Stanford University and is connected remotely to terminals in the schools by multiplexed telephone lines. Although the terminal has no graphic capability, it is a sturdy, low-cost device that provides the student with a printed copy of his interaction with the instructional program.)

automatically transfers him to the point in the curriculum where he finished on the previous day.

Reading instruction can be divided into two basic tasks which have been referred to as *decoding* and *communication*. Decoding is the rapid, if not automatic, association of phonemes or phoneme groups with their respective graphic representations. Communication involves reading for meaning, aesthetic enjoyment, emphasis, and the like. Our CAI program provides instruction in both types of tasks, but focuses primarily on decoding. The program is divided into eight parts or strands. As indicated in Figure 2, entry into a strand is determined by the student's level of achievement in the other strands. Instruction begins in Strand O, which teaches the skills required to interact with the program. Entry into the other strands is dependent on the student's performance in earlier strands. For example, the letter identification strand starts with a subset of letters used in the earliest sight words. When a student reaches a point in the letter identification strand where he has exhibited mastery over the letters used in the first words of the sight-word strand, he enters that strand. Similarly, entry into the spelling pattern strand and the phonics strand is controlled by the student's placement in the sight-word strand. On any given day a student may be seeing exercises drawn from as many as five strands. The dotted vertical lines in Figure 2 represent *maximal rate contours*, which control the student's progress in each strand relative to his progress in other strands. The rationale underlying these contours is that learning particular material in one strand facilitates learning in another strand; thus, the contours are constructed so that the student learns specific items from one strand in conjunction with specific items from other strands. In general, a student receives an amount of time in each strand proportional to the number of items yet to be completed in that strand before he reaches the next contour.

The CAI program is highly individualized so that a trace through the curriculum is unique for each student. The problem confronting the psychologist is to specify how a given subject's response history should be used to make instructional decisions. The approach that we have adopted is to develop simple mathematical models for the acquisition of the various skills in the cur-

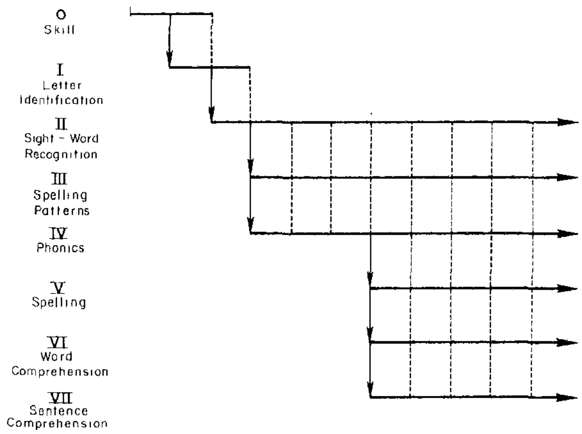

Figure 2. Schematic presentation of the strand structure. (Entry into each strand depends on a student's performance in earlier strands. The vertical dotted lines represent maximal rate contours which control the student's progress in each strand relative to the other strands.)

riculum and then use these models to specify optimal sequencing schemes. Basically, this problem is what has come to be known in the mathematical and engineering literature as *optimal control theory* or, more simply, *control theory*. The development of control theory has progressed at a phenomenal rate in the last decade, but most of the applications involve engineering or economic systems of one type or another. Precisely the same problems are posed in the area of instruction, except that the system to be controlled is the human learner rather than a machine or group of industries. If a model for the acquisition of a skill can be specified, then methods of control theory can be used to derive optimal instructional strategies.

I want to review some of the optimization procedures that have been developed, but in order for the reader to have some idea of how the CAI program operates, let me first describe a few of the simpler exercises used in Strands II, III, and IV.[4] Strand II provides for the development of a sight-word vocabulary. Vocabulary items are presented in five exercise formats, but only the copy exercise and the recognition exercise will be described here. The top panel of Table 1 illustrates the copy exercise, and the lower panel illustrates the recognition exercise. Note in Table 1 that when a student makes an error, the system responds with an audio message and prints out the correct response. In earlier versions of the program, the student was required to copy the correct response following an error. Experiments, however, demonstrated that the overt correction procedure was not effective and slowed down the pace of instruction; simply displaying the correct word following an error seems to be maximally effective.

Strand III provides practice with spelling patterns and emphasizes the regular grapheme–phoneme correspondences that exist in English. Table 2 illustrates exercises from this strand. For the exercise in the top panel of Table 2, the student is presented with three words involving the same spelling pattern and is required to select the correct one based on its initial letters. Once the student has learned to use the initial letter or letter sequence to distinguish between words, he then moves to the recall exercise illustrated in the bottom panel of Table 2. Here he works with a group of words, all involving the same spelling pattern. On each trial the audio system requests a word that requires adding an initial consonant or consonant cluster to the spelling pattern mastered in the preceding exercise. Whenever a student makes a correct response, a "+" sign is

TABLE 1

Examples of Two Exercises Used in Strand II (Sight-Word Recognition)

	Teletypewriter display	Audio message
	Copy exercise	
The program outputs:	PEN	(Type pen.)
The student responds by typing:	PEN	
The program outputs:	+	(Great!)
The program outputs:	EGG	(Type egg.)
The student responds by typing:	EFF	
The program outputs:	////EGG	(No, egg.)
	Recognition exercise	
The program outputs:	PEN NET EGG	(Type pen.)
The student responds by typing:	PEN	
The program outputs:	+	
The program outputs:	PEN EGG NET	(Type net.)
The student responds by typing:	NET	
The program outputs:	+	(Fabulous!)

Note. The top panel displays the copy exercise and the bottom panel the recognition exercise. Rows in the table correspond to successive lines on the teletypewriter printout.

[4] A detailed account of the curriculum can be found in Atkinson, Fletcher, Lindsay, Campbell, and Barr (1973).

TABLE 2

Examples of the Recognition and Recall Exercises Used in Strand III (Spelling Patterns)

	Teletypewriter display	Audio message
Recognition exercise		
The program outputs:	KEPT SLEPT CREPT	(Type kept.)
The student responds by typing:	KEPT	
The program outputs:	+	
Recall exercise		
The program outputs:		(Type crept.)
The student responds by typing:	CREPT	
The program outputs:	+	(That's fabulous!)

printed on the teletypewriter. In addition, every so often the program will give an audio feedback message; these messages vary from simple ones like "great," "that's fabulous," "you're doing brilliantly," to some that have cheering, clapping, or bells ringing in the background. These messages are not generated at random, but depend on the student's performance on that particular day. If his performance is above that of the preceding three days, it will be so recognized with frequent audio messages.

When the student has mastered a specified number of words in the sight-word strand, he begins exercises in the phonics strand; this strand concentrates on initial and final consonants and consonant clusters in combination with medial vowels. As in most linguistically oriented curricula, students are not required to rehearse or identify consonant sounds in isolation. The emphasis is on patterns of vowels and consonants that bear regular correspondences to phonemes. The phonic strand is the most complicated one of the group and involves eight exercise formats; only two of the formats will be described here. The upper panel of Table 3 illustrates an exercise in which the student is required to identify the graphic representation of phonemes occurring at the end of words. Each trial begins with an audio presentation of a word that includes the phonemes, and the student is asked to identify the graphic representation. After mastering this exercise, he is then transferred to the exercise illustrated in the bottom panel of Table 3. The same phonemes are presented, but now the student is requested to construct words by adding appropriate consonants.

Optimizing the Instructional Process

This has been a brief overview of some of the exercises used in the curriculum. The key to the curriculum is in the optimization schemes that control the sequencing of these exercises; these schemes can be classified at three levels. One level involves decision making within each strand. The problem is to decide which items to present for study, which exercise formats to present them in, and when to schedule review. A complete response history exists for each student, and this history is used to make trial-by-trial decisions regarding what instruction to present next. The second level of optimization deals with decisions about allocation of instruction time among the various strands for a given student. At the end of an instructional session, the student will have reached a certain point in each strand and a decision must be made as to the distribution of time among the strands in the next session. The third level of optimization deals with the distribution of instructional

TABLE 3

Examples of Two Exercises from Strand IV (Phonics)

	Teletypewriter display	Audio message
Recognition exercise		
The program outputs:	–IN –IT –IG	(Type /IG/ as in fig.)
The student responds by typing:	IG	
The program outputs:	+	(Good!)
The program outputs:	–IT –IN –IG	(Type /IT/ as in fit.)
The student responds by typing:	IT	
The program outputs:	+	
Build-a-word exercise		
The program outputs:	–IN –IT –IG P– –	(Type pin.)
The student responds by typing:	PIN	
The program outputs:	+	(Great!)
The program outputs:	–IG –IN –IT F– –	(Type fig.)
The student responds by typing:	FIN	
The program outputs:	////FIG	(No, we wanted fig.)

time among students. The question here is to allocate computer time among students to achieve instructional objectives that are defined not for the individual student but for the class as a whole. In some global sense, these three levels of optimization should be integrated into a unified program. However, our understanding of these matters is still very incomplete, and we have been satisfied to work with each in isolation, hoping that later they can be incorporated into a single package.

I want to describe some aspects of optimization on the first level and then go into detail on an example from the third level. Problems at the second level were touched on earlier when I described the strand structure of the curriculum and the use of maximum-rate contours to allocate time among strands. In some respects, this optimization program is the most interesting of the group, but it cannot be explained without going into considerable detail.[5]

Optimization within a strand, what has been called Level 1, can be illustrated using the sight-word strand. The strand comprises a list of about 1,000 words; the words are ordered in terms of their frequency in the student's vocabulary, and words at the beginning of the list tend to have highly regular grapheme–phoneme correspondences. At any point in time, a student will be working on a limited pool of words from the list; the size of this working pool depends on the student's ability level and is usually between 5 and 10 words. When one of these words is mastered, it is deleted from the pool and replaced by the next word on the list or by a word due for review. Figure 3 presents a flow chart for the strand. Each word in the working pool is in one of five possible instructional states. A trial involves sampling a word from the working pool and presenting it in an appropriate exercise format. The student is pretested on a word the first few times it is presented in order to eliminate words already known. If he knows the word, he will pass criterion for the pretest and it will be dropped from the working pool. If the student does not pass the pretest, he first studies the word in a series of trials, using the copy exercise, and then in a series of trials using the recognition exercise. If review is required, he studies the word again in what is designated in Figure 3 as Exercises 4 and 5.

[5] See Chant and Atkinson (1973) for a discussion of the problem and applications.

As indicated in the figure, a given word passes from one state to the next when it reaches criterion. And this presents the crux of the optimization problem, which is to define an appropriate criterion for each exercise. This has been done using simple mathematical models to describe the acquisition process for each exercise and the transfer functions that hold between exercises.[6] Basically, these models are simple Markov processes that have been extensively investigated by learning theorists and are known to provide reasonably accurate accounts of performance on our tasks. Parameters of the models are defined as functions of two factors: (*a*) the ability of the particular student and (*b*) the difficulty of the particular word. An estimate of the student's ability is obtained by analyzing his response record on all previous words, and an estimate of a word's difficulty is obtained by analyzing performance on that particular word for all students ever run on the program. The student records are continually updated by the computer and are used whenever necessary to compute a maximum likelihood estimate of each student's ability factor and each word's difficulty factor.[7] Given a well-defined model and estimates of its parameters, we can use the methods of control theory to define an optimal criterion for each exercise. The criterion will vary depending on the difficulty of the item, the student's ability level, and the precise sequence of correct and incorrect responses made by the student to the item. It is important to realize that the optimization scheme is not a simple one-stage branching program based on the last response, but rather depends, in a complicated way, on the student's complete response history.

Optimizing Class Performance

Now to turn to an example of optimization at what has been called Level 3. The effectiveness of the CAI program can be increased by optimally allocating instructional time among students. Suppose, for example, that a school has budgeted a fixed amount of time for CAI and must decide how to allocate that time among a class of first-grade students. For this example, maximizing the effec-

[6] For a discussion of the learning models, see Atkinson and Paulson (1972); for a discussion of the transfer models, see Dear and Atkinson (1962).

[7] See Atkinson and Paulson (1972) and Laubsch (1970) for a discussion of these methods.

tiveness of the CAI program is interpreted as meaning that we want to maximize student performance on a standardized reading test administered at the end of the first grade. Although other dependent measures can be used, this one provides a convenient benchmark against which to judge effectiveness.

On the basis of experimental data, the following equation has been derived that predicts performance on a standardized reading test as a function of the amount of time that a student spends on the CAI system:

$$P_i(t) = \alpha_i - (\alpha_i - \beta_i)e^{-\gamma_i t}.$$

The equation predicts Student i's performance on a standardized test as a function of the time, t, spent on the CAI system during the school year. The more time spent on the CAI program, the higher the level of achievement. The parameters α_i, β_i, and γ_i characterize Student i, and vary from one student to another. The parameters α_i and β_i are measures of Student i's maximal and minimal levels of achievement, respectively, and γ_i is a rate of progress measure. These parameters can be estimated from scores on reading readiness tests and from the student's performance during his first hour of CAI. After estimates of these parameters have been made, the above equation can be used to predict end-of-the-year test scores as a function of the CAI time allocated to that student.

To summarize, it is assumed that a school has budgeted a fixed amount of time T on the CAI system for a first-grade class of N students; further, students have had reading readiness tests and a preliminary run on the CAI system so that estimates of the parameters α, β, and γ can be made for each student. The problem then is to allocate time T among the N students so as to optimize learning. In order to do this, it is first necessary to have a model of the learning process. Although the above equation does not offer a very detailed account of learning, it suffices as a model for purposes of this problem, giving all the information that is required. This is an important point to keep in mind. The nature of the specific optimization problem determines the level of complexity that needs to be represented in the learning model. For some optimization problems, the model must provide a relatively complete account of learning in order to specify a viable strategy, but for other problems a simple descriptive equation often will suffice.

In addition to a model of the learning process, we must also specify our instructional objective. There are several objectives that seem reasonable, but only the following will be considered here:

A. Maximize the mean value of P over the class of students.

B. Minimize the variance of P over the class of students.

C. Maximize the mean value of P under the constraint that the resulting variance of P is less than or equal to the variance that would be obtained if no CAI were administered.

Objective A maximizes the gain for the class as a whole; Objective B reduces differences among students by making the class as homogeneous as possible; and Objective C attempts to maximize performance of the whole class, while insuring that differences among students are not amplified by CAI. To start, we will select Objective A as the instructional objective. If t_i is the time allocated to Student i, then the problem of deriving an optimal strategy reduces to maximizing the function

$$\phi(t_1,t_2,\cdots,t_N) = \frac{1}{N}\sum_{i=1}^{N}\left[\alpha_i - (\alpha_i - \beta_i)e^{-\gamma_i t_i}\right]$$

subject to the constraint that $t_1 + t_2 + \ldots + t_N = T$. This maximization can be done using the methods of dynamic programming. To illustrate the approach, computations were made for a first-grade class for which the parameters α, β, and γ had been estimated for each student. Employing these estimates, computations were carried out to determine the time allocations that maximized the above equation. For the optimal policy, the predicted mean performance level of the class on the end-of-the-year tests was 14% higher than a policy that allocated time equally among the students (i.e., a policy where $t_i = T/N$ for all students). This gain represents a substantial improvement; the drawback is that the class variance is roughly 15% greater than the variance for the class using an equal time policy. This means that if we are only interested in raising the class average, we

Figure 3. Partial flow chart for Strand II (sight-word recognition). (The various decisions represented in the bottom part of the chart are based on fairly complicated computations that make use of the student's response history.) →

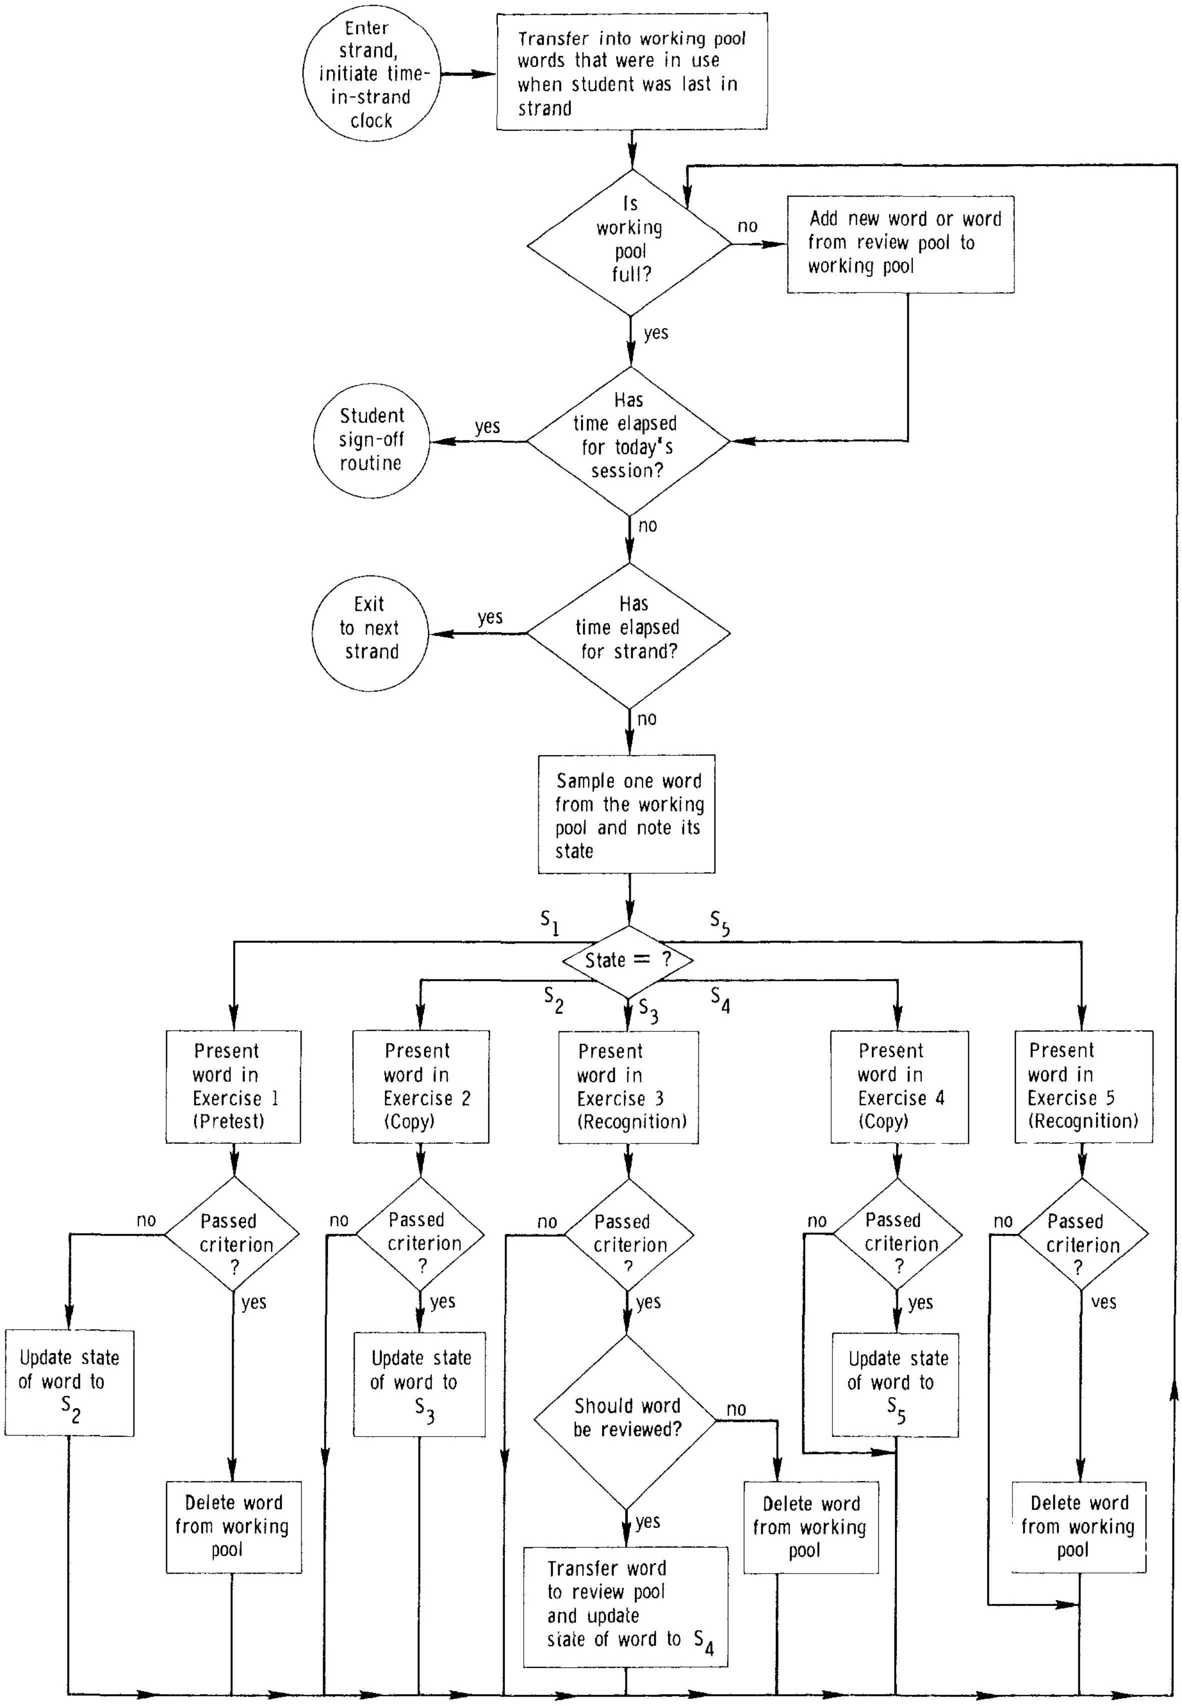
Enter strand, initiate time-in-strand clock
Transfer into working pool words that were in use when student was last in strand
Is working pool full?
no
Add new word or word from review pool to working pool
yes
Has time elapsed for today's session?
yes
Student sign-off routine
no
Has time elapsed for strand?
yes
Exit to next strand
no
Sample one word from the working pool and note its state
State = ?
S1
S5
S2
S3
S4
Present word in Exercise 1 (Pretest)
Present word in Exercise 2 (Copy)
Present word in Exercise 3 (Recognition)
Present word in Exercise 4 (Copy)
Present word in Exercise 5 (Recognition)
Passed criterion ?
no
yes
Update state of word to S2
Delete word from working pool
Update state of word to S3
Should word be reviewed?
Transfer word to review pool and update state of word to S4
Update state of word to S5
ves

TABLE 4

Percent Gain in the Mean of P and in the Variance of P When Compared with an Equal-Time Policy

	Instructional objective		
	A	B	C
% gain in mean of P	14	−15	8
% gain in variance of P	15	−12	−6

Note. P is an end-of-the-year performance score on a standardized reading test. In general, a policy that leads to a positive gain in the mean and a negative gain in variance is preferable.

will have to give the rapid learners substantial time on the CAI system and let them progress far beyond the slow learners.

Although a time allocation that complies with Objective A did increase overall class performance, the correlated increase in variance suggests that other objectives need to be considered. For comparison, time allocations also were computed for Objectives B and C. Table 4 presents the predicted gain in average class performance as a percentage of the mean value for the equal time policy. Objective B yielded a negative gain, and so it should, since its goal was to reduce variability, which is accomplished by holding back on rapid learners and giving a lot of attention to the slower ones. The reduction in variability for Objective B is 12%. Objective C, which strikes a balance between Objective A, on the one hand, and Objective B, on the other, yields an 8% gain in average performance and yet reduces variability by 6%.

In view of these results, Objective C seems to be the preferred one. It offers a substantial increase in average class performance while maintaining a low level of variability. These computations make it clear that the selection of an instructional objective should not be done in isolation but should involve a comparative analysis of several objectives, taking into account more than one dimension of performance. Even if the principal goal is to maximize the class average, it is inappropriate in most educational situations to select Objective A over C if it is only slightly better for the class average, while permitting variability to mushroom.

Effectiveness and Costs

Next, I want to make a few remarks about evaluation studies and the costs of the CAI program. Several evaluation studies have been conducted in the last few years, and yet another now is being done by the Educational Testing Service. Rather than merely review these studies, I would like to describe one in some detail.[8] In this particular study, 50 pairs of kindergarten students were matched on a number of variables, including sex and reading readiness scores. At the start of the first grade, one member of each pair was assigned to the experimental group and the other to the control group. Students in the experimental group received CAI, but only during the first grade; students in the control group received no CAI. The CAI lasted approximately 15 minutes per day; during this period the control group studied reading in the classroom. Except for this 15-minute period, the school day for the CAI group was like that of the control group. Standardized tests were administered at the end of the first grade and again at the end of the second grade. All the tests showed roughly the same pattern of results; to summarize the findings, only data from the California Cooperative Primary Reading Test will be described. At the end of the first grade, the experimental group showed a 5.05-month gain over the control group. The groups, when tested a year later (with no intervening CAI treatment), showed a difference of 4.90 months. Thus, the initial difference observed following one year of CAI was maintained, although not amplified, during the second year when no CAI was administered.

An interesting aspect of these results is that the boys appeared to benefit more from CAI than the girls. On all reading tests used in the evaluation, the girls as a group were superior to the boys. However, for the control group, the magnitude of the difference between boys and girls was greater than for the experimental group. For example, on the California Cooperative Primary Reading Test, the relative improvement for boys in the experimental group versus those in the control group was 42%; the corresponding figure for girls was 17%. These data suggest that both boys and girls benefit from CAI instruction, but that the gain is greater for boys. This observation is not unique to this study, but replicates a result from one of our earlier studies.[9] It is a common finding in

[8] For a detailed account of this study, see Fletcher and Atkinson (1972).

[9] See Atkinson (1968a) for a review of earlier studies.

research on reading that girls perform better than boys. This sex difference has been attributed at least in part to the social organization of the classroom and the predominance of female teachers. It is also argued that first-grade girls are more facile in memorization than boys and that this ability further aids the girls in curricula that emphasize sight-word procedures. If these two arguments are correct, then one would expect that placing the students in a CAI environment and using a curriculum that emphasizes analytic skills would reduce sex differences in reading. Our data tend to support this line of argument.

The results obtained in this and other studies can be used to project performance through the third grade. I will not explain the formula used to make the projection but will simply state the conclusion. For the population of students with which we have worked, the average reading level at the end of the third grade is approximately 2.9 when CAI is not used. For students who receive CAI, the grade level is 4.1. These values are to be compared with a national norm of 4.0. Thus, students with CAI are slightly above grade level by the end of the third grade, while those without CAI are about one year behind.

Can CAI be cost effective? The cost of daily sessions on our system is about $.55 per student. Based on a school year of 176 days, the yearly cost is roughly $97.00 per student. If this is multiplied by three, we have a figure of $291.00, a cost that places students at grade level by the end of the third grade who would normally be over a year behind. There is no doubt that such a cost is acceptable if future evaluations are as promising as they have been to date.[10]

Conclusion

In my view, individualizing instruction is the key factor in successfully teaching reading. This does not mean that all phases of instruction should be individualized, but often certain skills can be mastered only if instruction is sensitive to the student's particular difficulties. A teacher interacting on a one-to-one basis with a student may well be more effective than a CAI program. However, when working with a group of children (even as few as four or five), it is unlikely that she can match the computer's effectiveness over an extended period of time.

The possibilities for developing optimal instructional procedures are, of course, most promising at the elementary school level. In areas like initial reading and primary grade mathematics, we have an adequate understanding of many of the psychological processes involved. Simple models can be formulated to describe these processes, and in turn be used to derive optimal procedures. However, we know very little about the cognitive processes that underlie mastering a college-level curriculum in fields such as sociology or philosophy. In these cases, we cannot formulate models that describe learning, and thus cannot use the methods of optimization discussed here. However, more can be done at the college level than one might expect based on current work. Certainly, models can be developed for some aspects of learning in the natural sciences and second-language acquisition, even if these models are little more than descriptive equations. Research needs to be done to determine the feasibility of the approach, and much can be learned by experimenting with alternative optimization schemes, however loosely they may be related to formal models. Finding optimization schemes that work can tell us about the nature of the learning process and provide direction for theoretical analysis. A two-way exchange between the formulation of optimization procedures and development of descriptive models has not played as significant a role in psychological research as it merits.

In this discussion, I have tried to indicate some of the issues that arise in constructing a CAI program. From a broader perspective, the problem that we face in developing curricula is twofold. One aspect of the problem is to invent effective exercises for teaching particular skills and concepts; the other is to devise schemes for sequencing among these exercises. The problem of inventing and evaluating instructional exercises is an old one in psychology. There is no question that we know how to carry out the kinds of experiments needed for establishing the effectiveness of an instructional procedure. But the problem of formulating a scheme for sequencing among instructional procedures that is sensitive to the student's current state of knowledge is another matter, and one that has received very little investigation. Before the advent of the computer there was no real flexibility in manipulating the flow of instruction in a school situation; therefore, whether

[10] For a more detailed discussion of cost effectiveness, see Jamison, Fletcher, Suppes, and Atkinson (1974).

or not we understood how to individualize learning was of limited consequence. Now, with the computer, a new dimension of school learning has emerged. It is my belief that psychology's potential contributions in this area are of great practical significance. The development of a viable theory of instruction may be the most important issue facing psychology, and one that can revolutionize our conceptions of how man thinks and learns.

REFERENCES

Atkinson, R. C. Computerized instruction and the learning process. *American Psychologist,* 1968, **23**, 225–239. (a)

Atkinson, R. C. Learning to read under computer control. *Programmed Learning and Educational Technology: British Journal of the Association for Programmed Learning,* 1968, **5**, 25–37. (b)

Atkinson, R. C. Computer-assisted learning in action. *Proceedings of the National Academy of Sciences,* 1969, **63**, 588–594.

Atkinson, R. C. Ingredients for a theory of instruction. *American Psychologist,* 1972, **27**, 921–931.

Atkinson, R. C., Fletcher, J. D., Lindsay, E. J., Campbell, J. O., & Barr, A. Computer-assisted instruction in initial reading. *Educational Technology,* 1973, **13**, 27–37.

Atkinson, R. C., & Paulson, J. A. An approach to the psychology of instruction. *Psychological Bulletin,* 1972, **78**, 49–61.

Chant, V. G., & Atkinson, R. C. Optimal allocation of instructional effort to interrelated learning strands. *Journal of Mathematical Psychology,* 1973, **10**, 1–25.

Dear, R. E., & Atkinson, R. C. Optimal allocation of items in a simple, two-concept automated teaching model. In J. E. Coulson (Ed.), *Programmed learning and computer-based instruction.* New York: Wiley, 1962.

Fletcher, J. D., & Atkinson, R. C. Evaluation of the Stanford CAI program in initial reading. *Journal of Educational Psychology,* 1972, **63**, 597–602.

Groen, G. J., & Atkinson, R. C. Models for optimizing the learning process. *Psychological Bulletin,* 1966, **66**, 309–320.

Jamison, D., Fletcher, J. D., Suppes, P., & Atkinson, R. C. Cost and performance of computer-assisted instruction for compensatory education. In R. Radner & J. Froomkin (Eds.), *Education as an industry.* New York: Columbia University Press, 1974, in press.

Laubsch, J. H. Optimal item allocation in computer-assisted instruction. *IAG Journal,* 1970, **3**, 295–311.