11 INGREDIENTS FOR A THEORY OF INSTRUCTION [1]

RICHARD C. ATKINSON [2]

Stanford University

THE term "theory of instruction" has been in widespread use for over a decade and during that time has acquired a fairly specific meaning. By consensus it denotes a body of theory concerned with optimizing the learning process; stated otherwise, the goal of a theory of instruction is to prescribe the most effective methods for acquiring new information, whether in the form of higher order concepts or rote facts. Although usage of the term is widespread, there is no agreement on the requirements for a theory of instruction. The literature provides an array of examples ranging from speculative accounts of how children should be taught in the classroom to formal mathematical models specifying precise branching procedures in computer-controlled instruction.[3] Such diversity is healthy; to focus on only one approach would not be productive in the long run. I prefer to use the term theory of instruction to encompass both experimental and theoretical research, with the theoretical work ranging from general speculative accounts to specific quantitative models.

The literature on instructional theory is growing at a rapid rate. So much so that, at this point, a significant contribution could be made by someone willing to write a book summarizing and evaluating work in the area. I am reminded here of Hilgard's (1948) book, *Theories of Learning;* it played an important role in the development of learning theory by effectively summarizing alternative approaches and placing them in perspective. A book of this type is needed now in the area of instruction. The present article provides an overview of one of the chapters that I would like to see included in such a book; a title for the chapter might be "A Decision-Theoretic Analysis of Instruction." Basically, I consider here the factors that need to be examined in deriving optimal instructional strategies, and then I use this analysis to identify the key elements of a theory of instruction.

A Decision-Theoretic Analysis of Instruction

The derivation of an optimal strategy requires that the instructional problem be stated in a form amenable to a decision-theoretic analysis. Analyses based on decision theory vary somewhat from field to field, but the same formal elements can be found in most of them. As a starting point, I think it useful to identify these elements in a general way, and then to relate them to an instructional situation. They are as follows:

1. The possible states of nature.
2. The actions that the decision maker can take to transform the state of nature.
3. The transformation of the state of nature that results from each action.
4. The cost of each action.
5. The return resulting from each state of nature.

In the context of instruction, these elements divide naturally into three groups. Elements 1 and 3 are concerned with a description of the learning process; Elements 4 and 5 specify the cost-benefit dimensions of the problem; and Element 2 requires that the instructional actions from which the decision maker is free to choose be specified precisely.

[1] A briefer version of this article was presented as an invited address at the meetings of the American Educational Research Association, Chicago, April 1972. This research was sponsored in part by National Science Foundation Grant No. NSF GJ-443X2 and by Office of Naval Research Contract No. N00014-67-A-0112-0054.

[2] Requests for reprints should be sent to Richard C. Atkinson, Department of Psychology, Stanford University, Stanford, California 94305.

[3] See, for example, Smallwood (1962), Carroll (1963), Hilgard (1964), Bruner (1966), Groen and Atkinson (1966), Crothers and Suppes (1967), Gagné (1970), Seidel and Hunter (1970), Pask and Scott (1971), and Atkinson and Paulson (1972).

For the decision problems that arise in instruction, Elements 1 and 3 require that a model of the learning process exist. It is usually natural to identify the states of nature with the learning states of the student. Specifying the transformation of the states of nature caused by the actions of the decision maker is tantamount to constructing a model of learning for the situation under consideration. The learning model will be probabilistic to the extent that the state of learning is imperfectly observable or the transformation of the state of learning that a given instructional action will cause is not completely predictable.

The specification of costs and returns in an instructional situation (Elements 4 and 5) tends to be straightforward when examined on a short-term basis, but virtually intractable over the long term. For the short term, one can assign costs and returns for the mastery of, say, certain basic reading skills, but sophisticated determinations for the long-term value of these skills to the individual and society are difficult to make. There is an important role for detailed economic analyses of the long-term impact of education, but such studies deal with issues at a more global level than are considered here. The present analysis is limited to those costs and returns directly related to a specific instructional task.

Element 2 is critical in determining the effectiveness of a decision-theory analysis; the nature of this element can be indicated by an example. Suppose one wants to design a supplementary set of exercises for an initial reading program that involve both sight-word identification and phonics. Assume that two exercise formats have been developed, one for training on sight words, the other for phonics. Given these formats, there are many ways to design an overall program. A variety of optimization problems can be generated by fixing some features of the curriculum and leaving others to be determined in a theoretically optimal manner. For example, it may be desirable to determine how the time available for instruction should be divided between phonics and sight-word recognition, with all other features of the curriculum fixed. A more complicated question would be to determine the optimal ordering of the two types of exercises in addition to the optimal allocation of time. It would be easy to continue generating different optimization problems in this manner. The main point is that varying the set of actions from which the decision maker is free to choose changes the decision problem, even though the other elements remain the same.

Once these five elements have been specified, the next task is to derive the optimal strategy for the learning model that best describes the situation. If more than one learning model seems reasonable a priori, then competing candidates for the optimal strategy can be deduced. When these tasks have been accomplished, an experiment can be designed to determine which strategy is best. There are several possible directions in which to proceed after the initial comparison of strategies, depending on the results of the experiment. If none of the supposedly optimal strategies produces satisfactory results, then further experimental analysis of the assumptions of the underlying learning models is indicated. New issues may arise even if one of the procedures is successful. In the second example that I discuss, the successful strategy produces an unusually high error rate during learning, which is contrary to a widely accepted principle of programmed instruction (Skinner, 1968). When anomalies such as this occur, they suggest new lines of experimental inquiry, and often require a reformulation of the learning model.[4]

Criteria for a Theory of Instruction

The discussion to this point can be summarized by listing four criteria that must be satisfied prior to the derivation of an optimal instructional strategy:

1. A model of the learning process.
2. Specification of admissible instructional actions.
3. Specification of instructional objectives.
4. A measurement scale that permits costs to be assigned to each of the instructional actions and payoffs to the achievement of instructional objectives.

If these four elements can be given a precise interpretation, then it is generally possible to derive an optimal instructional policy. The solution for an optimal policy is not guaranteed, but in recent

[4] For a more extensive discussion of some of these points, see Atkinson and Paulson (1972), Calfee (1970), Dear, Silberman, Estavan, and Atkinson (1967), Laubsch (1970), and Smallwood (1971).

years some powerful tools have been developed for discovering optimal or near optimal procedures if they exist.

The four criteria just listed, taken in conjunction with methods for deriving optimal strategies, define either a model of instruction or a theory of instruction. Whether the term theory or model is used depends on the generality of the applications that can be made. Much of my own work has been concerned with the development of specific models for specific instructional tasks; hopefully, the collection of such models will provide the groundwork for a general theory of instruction.

In terms of the above criteria, it is clear that a model or theory of instruction is in fact a special case of what has come to be known in the mathematical and engineering literature as *optimal control theory* or, more simply, *control theory* (Kalman, Falb, & Arbib, 1969). The development of control theory has progressed at a rapid rate both in the United States and abroad, 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. To the extent that the above four elements can be formulated explicitly, methods of control theory can be used in deriving optimal instructional strategies.

To make some of these ideas more precise, I consider here two examples. One involves a *response-insensitive strategy* and the other a *response-sensitive strategy*. A response-insensitive strategy orders the instructional materials without taking into account the student's responses (except possibly to provide corrective feedback) as he progresses through the curriculum. In contrast, a response-sensitive strategy makes use of the student's response history in its stage-by-stage decisions regarding which curriculum materials to present next. Response-insensitive strategies are completely specified in advance and consequently do not require a system capable of branching during an instructional session. Response-sensitive strategies are more complex, but have the greatest promise for producing significant gains, for they must be at least as good, if not better, than the comparable response-insensitive strategy.

Optimizing Instruction in Initial Reading

The first example is based on work concerned with the development of a computer-assisted instruction (CAI) program for teaching reading in the primary grades (Atkinson & Fletcher, 1972). The program provides individualized instruction in reading and is used as a supplement to normal classroom teaching; a given student may spend anywhere from 0 to 30 minutes per day at a CAI terminal. For present purposes, only one set of results are considered, where the dependent measure is performance on a standardized reading achievement test administered at the end of the first grade. Using the Atkinson and Fletcher data, a statistical model can be formulated that predicts test performance as a function of the amount of time the student spends on the CAI system. Specifically, let $P_i(t)$ be student i's performance on a reading test administered at the end of first grade, given that he spends time t on the CAI system during the school year. Then within certain limits, the following equation holds:

$$P_i(t) = \alpha_i - \beta_i \exp(-\gamma_i t)$$

Depending on a student's particular parameter values, the more time spent on the CAI program, the higher the level of achievement at the end of the year. The parameters α, β, and γ characterize a given student and vary from one student to the next; α and $(\alpha - \beta)$ are measures of the student's maximal and minimal levels of achievement, respectively, and γ is a rate of progress measure. These parameters can be estimated from a student's response record obtained during his first hour of CAI. Stated otherwise, data from the first hour of CAI can be used to estimate the parameters α, β, and γ for a given student, and then the above equation enables one to predict end-of-year performance as a function of the CAI time allocated to that student.

The optimization problem that arises in this situation is as follows: Suppose that a school has budgeted a fixed amount of time T on the CAI system for the school year and must decide how to allocate the time among a class of n first-grade students. Assume, further, that all students have had a preliminary run on the CAI system so that estimates of the parameters α, β, and γ have been obtained for each student.

Let t_i be the time allocated to student i. Then the goal is to select a vector $(t_1, t_2, \ldots, t_n)$ that optimizes learning. To do this, one must check the four criteria for deriving an optimal strategy.

The first criterion is that there be a model of the learning process. The prediction equation for $P_i(t)$ does not offer a very complete account of learning; for purposes of this problem, however, the equation suffices as a model of the learning process, giving all of 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 must be represented in the learning model. For some problems, the model must provide a relatively complete account of learning in order to derive an optimal strategy, but for other problems a simple descriptive equation of the sort presented above will suffice.

The second criterion requires that the set of admissible instructional actions be specified. For the present case, the potential actions are simply all possible vectors $(t_1, t_2, \ldots, t_n)$ such that the t_is are nonnegative and sum to T. The only freedom decision makers have in this situation is in the allocation of CAI time to individual students.

The third criterion requires that the instructional objective be specified. There are several objectives that could be chosen in this situation. Consider four possibilities:

(*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 number of students who score at grade level at the end of the first year.
(*d*) Maximize the mean value of P satisfying the constraint that the resulting variance of P is less than or equal to the variance that would have been obtained if no CAI were administered.

Objective *a* maximizes the gain for the class as a whole; Objective *b* aims to reduce differences among students by making the class as homogeneous as possible; Objective *c* is concerned specifically with those students who fall behind grade level; Objective *d* attempts to maximize performance of the whole class but insures that differences among students are not amplified by CAI. Other instructional objectives can be listed, but these are the ones that seemed most relevant. For expository purposes, I have selected *a* as the instructional objective.

The fourth criterion requires that costs be assigned to each of the instructional actions and that payoffs be specified for the instructional objectives. In the present case, one can assume that the cost of CAI does not depend on how time is allocated among students and that the measurement of payoff is directly proportional to the students' achieved value of P.

In terms of the four criteria, the problem of deriving an optimal instructional strategy reduces to maximizing the function

$$\phi(t_1,t_2,\ldots,t_n) = \frac{1}{n}\sum_{i=1}^{n} P_i(t_i)$$

$$= \frac{1}{n}\sum_{i=1}^{n} \left[\alpha_i - \beta_i \exp(-\gamma_i t_i)\right]$$

subject to the constraint that

$$\sum_{i=1}^{n} t_i = T$$

and

$$t_i \geq 0.$$

This maximization can be done by using the method of dynamic programming (Bellman, 1961). In order 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, $\bar{P}$, was 15% higher than a policy that allocated time equally to students (i.e., a policy in which $t_i = t_j$ for all i and j). This gain represents a substantial improvement; the drawback is that the variance of the P scores is roughly 15% greater than for the equal-time policy. This means that if one were interested primarily in raising the class average, one would have to let the rapid learners move ahead and 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 leads one to believe that other objectives might be more beneficial. For comparison, time allocations also were computed for Objectives *b*, *c*, and *d*. Figure 1 presents the predicted gain in $\bar{P}$ as a percentage of $\bar{P}$ for the equal-time policy. Objectives *b* and *c* yield negative gains, and so they should since their goal is to reduce variability, which is accomplished by holding back on the rapid learners and giving a lot of attention to the slower ones. The reduction in variability for these two objectives, when compared with the equal-time policy, is 12% and 10%, respectively. Objective *d*, which attempts to strike a balance between Objective *a* on the one hand and Objectives *b* and *c* on the other, yields an 8% increase in $\bar{P}$ and yet reduces variability by 6%.

In view of these computations, Objective *d* seems to be preferred; it offers a substantial increase in mean performance while maintaining a low level of variability. As yet, this policy has not been implemented, so only theoretical results can be reported. Nevertheless, these examples yield differences that illustrate the usefulness of this type of analysis. They make it clear that the selection of an instructional objective should not be done in isolation, but should involve a comparative analysis of several alternatives taking into account more than one dimension of performance. For example, even if the principal goal is to maximize $\bar{P}$, it would be inappropriate in most educational situations to select a given objective over some other if it yielded only a small average gain while variability mushroomed.

Techniques of the sort presented above have been developed for other aspects of the CAI reading program. One of particular interest involves deciding for each student, on a week-by-week basis, how time should be divided between training in phonics and in sight-word identification (Chant & Atkinson, in press). However, these developments are not considered here; it is more useful to turn to another example of a quite different type.

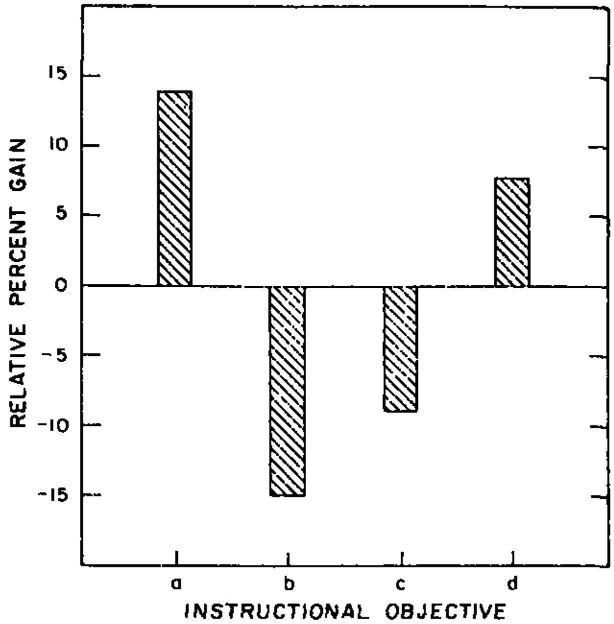

FIG. 1. Percentage gains in the mean value of *P* when compared with an equal-time policy for four policies each based on a different instructional objective.

Optimizing the Learning of a Second-Language Vocabulary

The second example deals with learning a foreign language vocabulary. A similar example could be given from our work in initial reading, but this particular example has the advantage of permitting us to introduce the concept of learner-controlled instruction. In developing the example, I consider first some experimental work comparing three instructional strategies and only later explain the derivation of the optimal strategy.[5]

The goal is to individualize instruction so that the learning of a second-language vocabulary occurs at a maximum rate. The constraints imposed on the task are typical of a school situation. A large set of German–English items are to be learned during an instructional session that involves a series of trials. On each trial one of the German words is presented, and the student attempts to give the English translation; the correct translation then is presented for a brief study period. A predetermined number of trials is allocated for the instructional session, and after an intervening period of one week a test is administered over the entire vocabulary. The optimization problem is to formulate a strategy for presenting items during the instructional session so that performance on the delayed test will be maximized.

Three strategies for sequencing the instructional material are considered here. One strategy (designated the random-order strategy) is simply to cycle through the set of items in a random order; this

[5] A detailed account of this research can be found in Atkinson (in press).

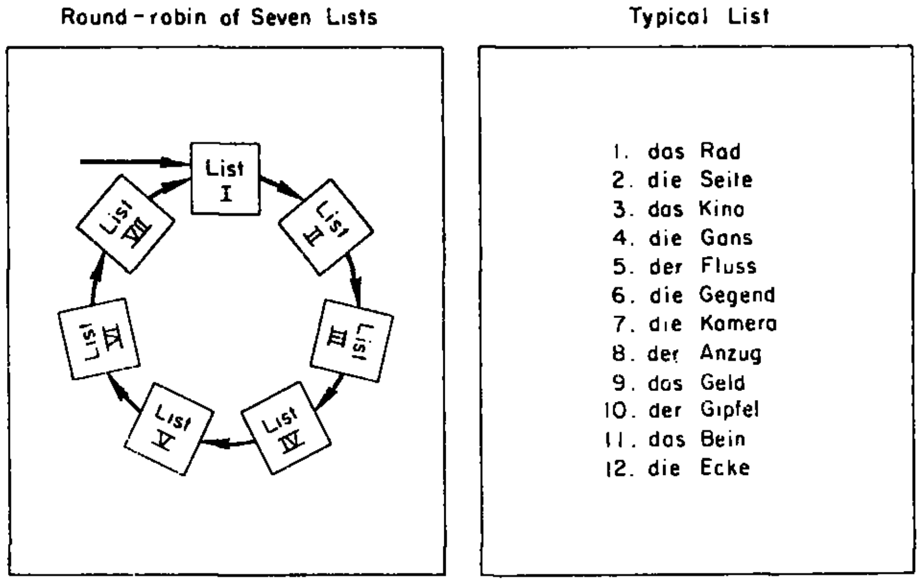

FIG. 2. Schematic representation of the round robin of display lists and an example of one such list.

strategy is not expected to be particularly effective, but it provides a benchmark against which to evaluate others. A second strategy (designated the learner-controlled strategy) is to let the student determine for himself how best to sequence the material. In this mode the student decides on each trial which item is to be tested and studied; the learner, rather than an external controller, determines the sequence of instruction. The third scheme (designated the response-sensitive strategy) is based on a decision-theoretic analysis of the instructional task. A *mathematical model of learning* that has provided an accurate account of vocabulary acquisition in other experiments is assumed to hold in the present situation. This model is used to compute, on a trial-by-trial basis, an individual student's current state of learning. Based on these computations, items are selected from trial to trial so as to optimize the level of learning *achieved at the termination of the instructional* session. The details of this strategy are complicated and can be discussed more meaningfully after the experimental procedure and results have been presented.

Instruction in this experiment is carried out under computer control. The students are required to participate in two sessions: an *instructional session* of approximately two hours and a briefer *delayed-test session* administered one week later. The delayed test is the same for all students and involves a test over the entire vocabulary. The instructional session is more complicated. The vocabulary items are divided into seven lists, each containing 12 German words; the lists are arranged in a round-robin order (see Figure 2). On

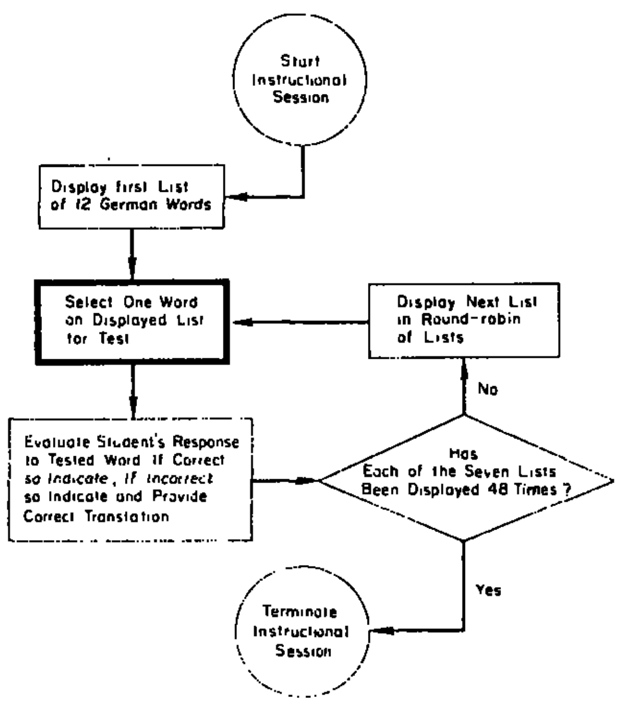

FIG. 3. Flow chart describing the trial sequence during the instructional session. The selection of a word for test on a given trial (box with heavy border) varied over experimental conditions.

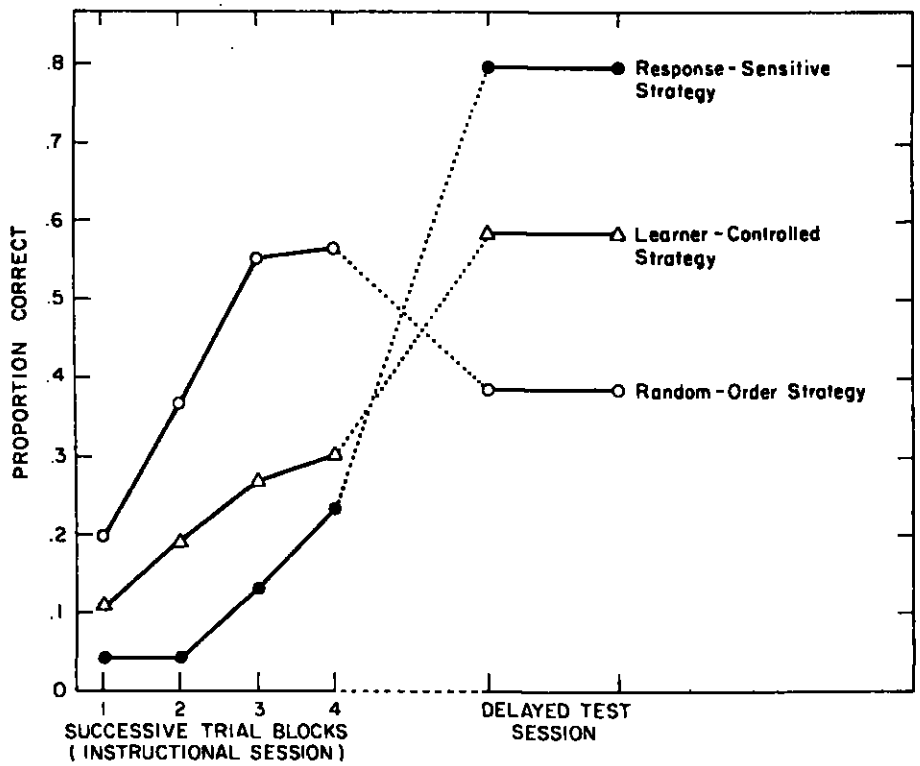

FIG. 4. Proportion of correct responses in successive trial blocks during the instructional session, and on the delayed test administered one week later.

each trial of the instructional session a list is displayed, and the student inspects it for a brief period of time. Then one of the items on the displayed list is selected for test and study. In the random-order and response-sensitive conditions, the item is selected by the computer. In the learner-controlled condition, the item is chosen by the student. After an item has been selected for test, the student attempts to provide a translation; then feedback regarding the correct translation is given. The next trial begins with the computer displaying the next list in the round robin, and the same procedure is repeated. The instructional session continues in this fashion for 336 trials (see Figure 3).

The results of the experiment are summarized in Figure 4. Data are presented on the left side of the figure for performance on successive blocks of trials during the instructional session; on the right side are results from the test session administered one week after the instructional session. Note that during the instructional session the probability of a correct response is highest for the random-order condition, next highest for the learner-controlled condition, and lowest for the response-sensitive condition. The results, however, are reversed on the delayed test. The response-sensitive condition is best by far with 79% correct; the learner-controlled condition is next with 58%; and the random-order condition is poorest at 38%. The observed pattern of results is expected. In the learner-controlled condition, the students are trying, during the instructional session, to test and study those items they do not know, and they should have a lower score than students in the random-order condition where testing is random and includes many items already mastered. The response-sensitive procedure also attempts to identify for test and study those items that have not yet been mastered and thus also produces a high error rate during the instructional session. The ordering of groups on the delayed test is reversed since now the entire set of words is tested; when all items are tested, the probability of a correct response tells how much of the list actually has been mastered. The magnitude of the effects observed on the delayed test is large and of practical significance.

Now that the effectiveness of the response-sensitive strategy has been established, I turn to a discussion of how it was derived. The strategy is based on a model of vocabulary learning that has been investigated in the laboratory and has been shown to be quite accurate (Atkinson, in press; Atkinson & Crothers, 1964). The model assumes that a given item is in one of three states (P, T, and U) at any moment in time. If the item is in State P, then its translation is known, and this knowledge is "relatively" permanent in the sense that the learning of other vocabulary items will not interfere with it. If the item is in State T, then it is also known but on a "temporary" basis; in State T, other items can give rise to interference effects that cause the item to be forgotten. In State U, the item is not known, and the student is unable to provide a translation. Thus, in States P and T a correct translation is given with probability 1, whereas in State U the probability is 0.

When a test and study occur on a given item, the following transition matrix describes the possible change in state from the onset of the trial to its termination:

$$
\begin{array}{cc}
 & \begin{array}{ccc} \mathrm{P} & \mathrm{T} & \mathrm{U} \end{array} \\
\mathbf{A} = \begin{array}{c} \mathrm{P} \\ \mathrm{T} \\ \mathrm{U} \end{array} & \begin{bmatrix} 1 & 0 & 0 \\ a & 1-a & 0 \\ bc & (1-b)c & 1-c \end{bmatrix}
\end{array}
$$

Rows of the matrix represent the state of the item at the start of the trial, and the columns represent its state at the end of the trial. On a trial when some other item is presented for test and study, a transition in the learning state of the original item also may take place; namely, forgetting is possible in the sense that if the item is in State T, it may transit into State U. This forgetting can occur only if the student makes an error on the other item; in that case the transition matrix applied to the original item is as follows:

$$
\begin{array}{cc}
 & \begin{array}{ccc} \mathrm{P} & \mathrm{T} & \mathrm{U} \end{array} \\
\mathbf{F} = \begin{array}{c} \mathrm{P} \\ \mathrm{T} \\ \mathrm{U} \end{array} & \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1-f & f \\ 0 & 0 & 1 \end{bmatrix}
\end{array}
$$

To summarize, consider the application of Matrices **A** and **F** to some specific item on the list; when the item itself is presented for test and study, transition Matrix **A** is applied; when some other item is presented that elicits an error, then Matrix **F** is applied. The above assumptions provide a complete account of the learning process. The parameters in Matrices **A** and **F** measure the difficulty level of a German–English pair and vary across items. On the basis of prior experiments, numerical estimates of these parameters exist for each of the items used in the experiment.

As noted earlier, the formulation of a strategy requires that one be precise about the quantity to be maximized. For the present task, the goal is to maximize the number of items correctly translated on the delayed test. To do this, a theoretical relationship must be specified between the state of learning at the end of the instructional session and performance on the delayed test. The assumption made here is that only those items in State P at the end of the instructional session will be translated correctly on the delayed test; an item in State T is presumed to be forgotten during the intervening week. Thus, the problem of maximizing delayed-test performance involves, at least in theory, maximizing the number of items in State P at the termination of the instructional session.

Having numerical values for parameters and knowing the student's response history, it is possible to estimate his current state of learning.[6] Stated more precisely, the learning model can be used to derive equations and, in turn, compute the probabilities of being in States P, T, and U for each item at the start of trial n, conditionalized on the student's response history up to and including trial $n-1$. Given numerical estimates of these probabilities, a strategy for optimizing performance is to select that item for presentation (from the current display list) that has the greatest probability of moving into State P if it is tested and studied on the trial. This strategy has been termed the one-stage optimization procedure because it looks ahead one trial in making decisions. The true optimal policy (i.e., an N-stage procedure) would consider all possible item-response sequences for the remaining trials and select the next item so as to maximize the number of items in State P at the termination of the instructional session. For the present case, the N-stage policy cannot be applied

[6] The student's response history is a record (for each trial) of the item presented and the response that occurred. It can be shown that a sufficient history exists which contains only the information necessary to estimate the student's current state of learning; the sufficient history is a function of the complete history and the assumed learning model. For the model considered here, the sufficient history is fairly simple but cannot be readily described without extensive notation.

because the necessary computations are too time consuming even for a large computer. Fortunately, Monte Carlo studies indicate that the one-stage policy is a good approximation to the optimal strategy for a variety of Markov learning models; it was for this reason, as well as the relative ease of computing, that the one-stage procedure was employed.[7] The computational procedure described above was implemented on the computer and permitted decisions to be made on-line for each student on a trial-by-trial basis.

The response-sensitive strategy undoubtedly can be improved on by elaborating the learning model. Those familiar with developments in learning theory will see a number of ways of introducing more complexity into the model and thereby increasing its precision. I do not pursue such considerations here, however, since my reason for presenting the example was not to theorize about the learning process but rather to demonstrate how a simple learning model can be used to define an instructional procedure.

Concluding Remarks

Hopefully, these two examples illustrate the steps involved in developing an optimal strategy for instruction. Both examples deal with relatively simple problems and thus do not indicate the range of developments that have been made or that are clearly possible. It would be a mistake, however, to conclude that this approach offers a solution to the problems facing education. There are some fundamental obstacles that limit the generality of the work.

The major obstacles may be identified in terms of the four criteria that were specified as prerequisites for an optimal strategy. The first criterion concerns the formulation of learning models. The models that now exist are totally inadequate to explain the subtle ways by which the human organism stores, processes, and retrieves information. Until there is a much deeper understanding of learning, the identification of truly effective strategies will not be possible. However, an all-inclusive theory of learning is not a prerequisite for the development of optimal procedures. What is needed instead is a model that captures the essential features of that part of the learning process being tapped by a given instructional task. Even models that may be rejected on the basis of laboratory investigation can be useful in deriving instructional strategies. The two learning models considered in this article are extremely simple, and yet the optimal strategies they generate are quite effective. My own preference is to formulate as complete a learning model as intuition and data will permit and then to use that model to investigate optimal procedures; when possible the learning model will be represented in the form of mathematical equations but otherwise as a set of statements in a computer-simulation program. The main point is that the development of a theory of *instruction cannot progress if one holds the view* that a complete theory of learning is a prerequisite. Rather, advances in learning theory will affect the development of a theory of instruction, and conversely the development of a theory of instruction will influence research on learning.

The second criterion for deriving an optimal strategy requires that admissible instructional *actions be specified* clearly. The set of potential instructional inputs places a definite limit on the effectiveness of the optimal strategy. In my opinion, powerful instructional strategies must necessarily be adaptive; that is, they must be sensitive on a moment-to-moment basis to a learner's unique response history. My judgment on this matter is based on limited experience, restricted primarily to research on teaching initial reading. In this area, however, the evidence seems to be absolutely clear: the manipulation of method variables accounts for only a small percentage of the variance when not accompanied by instructional strategies that permit individualization. Method variables like the modified teaching alphabet, oral reading, the linguistic approach, and others undoubtedly have beneficial effects. However, these effects are minimal in comparison to the impact that is possible when instruction is adaptive to the individual learner. Significant progress in dealing with the nation's problem of teaching reading will require individually prescribed programs, and sophisticated programs will necessitate some degree of computer intervention either in the form of CAI or computer-managed instruction. As a corollary to this point, it is evident from observations of students in our CAI reading program that the more effective the adaptive strategy the less important are extrinsic

[7] For a discussion of one-stage and N-stage policies and Monte Carlo studies comparing them, see Groen and Atkinson (1966), Calfee (1970), and Laubsch (1970).

motivators. Motivation is a variable in any form of learning, but when the instructional process is truly adaptive the student's progress is sufficient reward in its own right.

The third criterion for an optimal strategy deals with instructional objectives, and the fourth with cost–benefit measures. In the analyses presented here, it was tacitly assumed that the curriculum material being taught is sufficiently beneficial to justify allocating time to it. Further, in both examples the costs of instruction were assumed to be the same for all strategies. If the costs of instruction are equal for all strategies, they may be ignored and attention focused on the comparative benefits of the strategies. This is an important point because it greatly simplifies the analysis. If both costs and benefits are significant variables, then it is essential that both be estimated accurately. This is often difficult to do. When one of these quantities can be ignored, it suffices if the other can be assessed accurately enough to order the possible outcomes. As a rule, both costs and benefits must be weighed in the analysis, and frequently subtopics within a curriculum vary significantly in their importance. In some cases, whether or not a certain topic should be taught at all is the critical question. Smallwood (1971) has treated problems similar to the ones considered in this article in a way that includes some of these factors in the structure of costs and benefits.

My last remarks deal with the issue of learner-controlled instruction. One way to avoid the challenge and responsibility of developing a theory of instruction is to adopt the view that the learner is the best judge of what to study, when to study, and how to study. I am alarmed by the number of individuals who advocate this position despite a great deal of negative evidence. Do not misinterpret this remark. There obviously is a place for the learner's judgments in making instructional decisions. In several CAI programs that I have helped develop, the learner plays an important role in determining the path to be followed through the curriculum. However, using the learner's judgment as one of several items of information in making an instructional decision is quite different from proposing that the learner should have complete control. My data, and the data of others, indicate that the learner is not a particularly effective decision maker. Arguments against learner-controlled programs are unpopular in the present climate of opinion, but they need to be made so that one will not be seduced by the easy answer that a theory of instruction is not required because "who can be a better judge of what is best for the student than the student himself."

This article illustrates the steps involved in deriving an optimal strategy and their implications for a theory of instruction. I want to emphasize a point made at the outset—namely, that the approach is only one of many that needs to be pursued. Obviously the main obstacle is that adequate theories as yet do not exist for the learning processes that we most want to optimize. However, as the examples indicate, analyses based on highly simplified models can be useful in identifying problems and focusing research efforts. It seems clear that this type of research is a necessary component in a program designed to develop a general theory of instruction.

REFERENCES

Atkinson, R. C. Optimizing the learning of a second-language vocabulary. *Journal of Experimental Psychology*, in press.

Atkinson, R. C., & Crothers, E. J. A comparison of paired-associate learning models having different acquisition and retention axioms. *Journal of Mathematical Psychology*, 1964, **1**, 285–315.

Atkinson, R. C., & Fletcher, J. D. Teaching children to read with a computer. *The Reading Teacher*, 1972, **25**, 319–327.

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

Bellman, R. *Adaptive control processes.* Princeton: Princeton University Press, 1961.

Bruner, J. S. *Toward a theory of instruction.* Cambridge: Harvard University Press, 1966.

Calfee, R. C. The role of mathematical models in optimizing instruction. *Scientia: Revue Internationale de Synthèse Scientifique*, 1970, **105**, 1–25.

Carroll, J. B. A model of school learning. *Teachers College Record*, 1963, **64**, 723–733.

Chant, V. G., & Atkinson, R. C. Optimal allocation of instructional effort to interrelated learning strands. *Journal of Mathematical Psychology*, in press.

Crothers, E. J., & Suppes, P. *Experiments in second-language learning.* New York: Academic Press, 1967.

Dear, R. E., Silberman, H. F., Estavan, D. P., & Atkinson, R. C. An optimal strategy for the presentation of paired-associate items. *Behavioral Science*, 1967, **12**, 1–13.

Gagné, R. M. *The conditions of learning.* (2nd ed.) New York: Holt, Rinehart & Winston, 1970.

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

HILGARD, E. R. *Theories of learning*. New York: Appleton-Century-Crofts, 1948.

HILGARD, E. R. (Ed.) *Theories of learning and instruction: The sixty-third yearbook of the National Society for the Study of Education*. Chicago: University of Chicago Press, 1964.

KALMAN, R. E., FALB, P. L., & ARBIB, M. A. *Topics in mathematical system theory*. New York: McGraw-Hill, 1969.

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

PASK, G., & SCOTT, B. C. E. Learning and teaching strategies in a transformational skill. *British Journal of Mathematical and Statistical Psychology*, 1971, **24**, 205–229.

SEIDEL, R. J., & HUNTER, H. G. The application of theoretical factors in teaching problem-solving by programed instruction. *International Review of Applied Psychology*, 1970, **19**, 41–81.

SKINNER, B. F. *The technology of teaching*. New York: Appleton-Century-Crofts, 1968.

SMALLWOOD, R. D. *A decision structure for teaching machines*. Cambridge: M.I.T. Press, 1962.

SMALLWOOD, R. D. The analysis of economic teaching strategies for a simple learning model. *Journal of Mathematical Psychology*, 1971, **8**, 285–301.