

UC Office of the President

Computer Assisted Instruction (CAI)

Title

An Approach to the Psychology of Instruction

Permalink

<https://escholarship.org/uc/item/7v73d6qs>

Authors

Atkinson, Richard C.

Paulson, J. A.

Publication Date

1972

Peer reviewed

AN APPROACH TO THE PSYCHOLOGY OF INSTRUCTION

R. C. ATKINSON¹ AND J. A. PAULSON

Stanford University

The relationship between a theory of learning and a theory of instruction is discussed. Examples are presented that illustrate how to proceed from a theoretical description of the learning process to the specification of an optimal strategy for carrying out instruction. The examples deal with fairly simple learning tasks and are admittedly of limited generality. Nevertheless, they clearly define the steps necessary for deriving and testing instructional strategies, thereby providing a set of procedures for analyzing more complex problems. The parameter-dependent optimization strategies are of particular importance because they take into account individual differences among learners as well as differences in difficulty among curriculum units. Experimental evaluations indicate that the parameter-dependent strategies lead to major gains in learning, when compared with strategies that do not take individual differences into account.

The task of relating the methods and findings of research in the behavioral sciences to the problems of education is a continuing concern of both psychologists and educators. A few years ago, when our faith in the ability of money and science to cure social ills was at its peak, an educational researcher could content himself with trying to answer the same questions that were being studied by his psychologist colleagues. The essential difference was that his studies referred explicitly to educational settings, whereas those undertaken by psychologists strived for greater theoretical generality. There was implicit confidence that as the body of behavioral research grew, applications to education would occur in the natural course of events. When these applications failed to materialize, confidence was shaken. Clearly, something essential was missing from educational research.

A number of factors contributed to the feeling that something was wrong with business as usual. Substantial curriculum changes, initiated on a national scale after the Soviet's launching of Sputnik, had to be carried out with only minimal guidance from behavioral scientists. Developers of programmed learning

and computer-assisted instruction faced similar problems. Although the literature in learning theory was perhaps more relevant to their concerns, the questions it treated were still not the critical ones from the viewpoint of instruction. This situation would not have been surprising had the study of learning been in its infancy. But far from that, the psychology of learning had a long and impressive history. An extensive body of experimental literature existed, and many simple learning processes were being described with surprising precision using mathematical models. Whatever was wrong, it did not seem to be a lack of scientific sophistication.

These issues were on the minds of those who contributed to the 1964 *Yearbook of the National Society for the Study of Education*, edited by Hilgard (1964). In that book Bruner (1964) summarized the feelings of many of the contributors when he called for a theory of instruction, which he sharply distinguished from a theory of learning. He emphasized that where the latter is essentially descriptive, the former should be prescriptive, setting forth rules specifying the most effective ways of achieving knowledge or mastering skills. This distinction served to highlight the difference in the goals of experiments designed to advance the two kinds of theory. In many instances variations in instructional procedures affect several psychological variables simultaneously. Experiments that are appropriate for comparing methods of instruction may be virtually impossible to interpret in terms of learning

¹ An earlier version of this paper was presented by the first author at a conference on "The Use of Computers in Education" organized by the Japanese Ministry of Education in collaboration with the Organization for Economic Cooperation and Development in Tokyo, July 1970.

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

theory because of this confounding of variables. The importance of developing a theory of instruction justifies experimental programs designed to explore alternative instructional procedures, even if the resulting experiments are difficult to place in a learning-theoretic framework.

The task of going from a description of the learning process to a prescription for optimizing learning must be clearly distinguished from the task of finding the appropriate theoretical description in the first place. However, there is a danger that preoccupation with finding prescriptions for instruction may cause us to overlook the critical interplay between the two enterprises. Developments in control theory (Bellman, 1961) and statistical decision theory (Raiffa & Schlaiffer, 1968) provide potentially powerful methods for discovering optimal decision-making strategies in a wide variety of contexts. In order to use these tools it is necessary to have a reasonable model of the process to be optimized. As noted earlier, some learning processes can already be described with the required degree of accuracy. This article examines an approach to the psychology of instruction which is appropriate when the learning is governed by such a process.

STEPS IN THE DEVELOPMENT OF OPTIMAL INSTRUCTIONAL STRATEGIES

The development of optimal strategies can be broken down into a number of tasks that involve both descriptive and normative analyses. One task requires that the instructional problem be stated in a form amenable to a decision-theoretic analysis. While the detailed formulations of decision problems vary widely from field to field, the same formal elements can be found in most of them. It will be a useful starting point to identify these elements in the context of an instructional situation.

The formal elements of a decision problem that must be specified are:

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.

Statistical aspects occur in a decision problem when uncertainty is associated with one or more of these elements. For example, the state of nature may be imperfectly observable or the transformation of the state of nature which a given action will cause may not be completely predictable.

In the context of the psychology of instruction, most of these elements divide naturally into two groups, those having to do with the description of the underlying learning process and those specifying the cost-benefit dimensions of the problem. The one element that does not fit is the specification of the set of actions from which the decision maker must make his choice. The nature of this element can be indicated by an example.

Suppose one wants to design a supplemental program of exercises for an initial reading program. Most reasonable programs of initial reading instruction include both training in sight-word identification and training in phonics. Let us assume that on the basis of experimentation two useful 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 design and leaving the 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 design 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 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.

For the decision problems that arise in instruction it is usually natural to identify the states of nature with 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 role of costs and returns is more formal than substantive for the class of decision problems considered in this article. The specification of costs and returns in instructional situations tends to be straightforward when examined on a short-time basis, but virtually intractable over the long term. In 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 analysis of the long-term impact of education, but such studies deal with issues at a more global level than we require. In this article analysis is limited to those costs and returns directly related to the specific instructional task being considered.

After a problem has been formulated in a way amenable to decision-theoretic analysis, the next step is to derive the optimal strategy for the learning model which 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 steps 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 one case that we discuss, the successful strategy produced an usually high error rate during learning, which is contrary to a widely accepted principle of programmed instruction. When anomalies such as this occur, they suggest new lines of experimental inquiry, and often require a reformulation of the axioms of the learning model. The learning model may have provided an excellent account of data for a range of experimental conditions but can prove totally inadequate in an optimization condition where

special features of the procedure magnify inaccuracies of the model that had previously gone undetected.

AN OPTIMIZATION PROBLEM THAT ARISES IN COMPUTER-ASSISTED INSTRUCTION

One application of computer-assisted instruction that has proved to be very effective in the primary grades involves a regular program of practice and review specifically designed to complement the efforts of the classroom teacher (Atkinson, 1969). Some of the curriculum materials in such programs take the form of lists of instructional units or items. The objective of the computer-assisted instruction programs is to teach students the correct response to each item in a given list. Typically, a sublist of items is presented each day in one or more fixed exercise formats. The optimization problem that arises concerns the selection of items for presentation on a given day.

The Stanford Reading Project is an example of such a program in initial reading instruction (Atkinson & Fletcher, 1972). The vocabularies of several of the commonly used basal readers were compiled into one dictionary, and a variety of exercises using these words were developed to teach reading skills. These exercises were designed principally to strengthen the student's decoding skills, with special emphasis on letter identification, sight-word recognition, phonics, spelling patterns, and word comprehension. The details of the teaching procedure vary from one exercise to another, but most include a sequence in which a curriculum item is presented, eliciting a response from the student, followed by a short period for studying the correct response. For example, one exercise in sight-word recognition has the following format:

Teletype Display	Audio Message
NUT MEN RED	Type red.

Three words are printed on the teletype, followed by an audio presentation of one of the words. Control is then turned over to the student; if he types the correct word a reinforcing message is given, and the computer program then proceeds to the next presentation. If the student responds incorrectly or exceeds the time, the teletype prints the

correct word simultaneously with its audio presentation and then moves to the next presentation. Under an early version of the program, items were presented in predetermined sublists, with an exercise continuing on a sublist until a specified criterion has been met.

Strategies can be found that will improve on the fixed order of presentation. Two studies described below are concerned with the development of such strategies. One study examines alternative presentation strategies for teaching spelling words to elementary school children, and the other examines strategies for teaching Swahili vocabulary items to college-level students. The optimization problems in both studies were essentially the same. A list of N items is to be learned, and a fixed number of days, D , are allocated for its study. On each day a sublist of items is presented for test and study. The sublist always involves M items, and each is presented only once for test followed by a study period. The total set of N items is extremely large with regard to any sublist of M items. Once the experimenter has specified a sublist for a given day, its order of presentation is random. After the D days of study are completed, a posttest is given over all items. The parameters N , D , and M are fixed, and so is the instructional format on each day. Within these constraints the problem is to maximize performance on a posttest by an appropriate selection of sublists from day to day. The strategy for selecting sublists from day to day is dynamic (or *response sensitive*, using the terminology of Groen & Atkinson, 1966) to the extent that it depends upon the student's prior history of performance.

Three Models of the Learning Process

Two extremely simple learning models are considered first. Then a third model, which combines features of the first two, is described.

In the first model, the state of the learner with respect to each item is completely determined by the number of times the item has been studied. At the start of the experiment an item has some initial probability of error; each *time* the item is presented its error probability is reduced by a factor α , which is less than one. Stated as a difference equation, the probability of an error on the $n + 1$ st presentation of an item is related to its probability on

the n th presentation as follows:

$$q_{n+1} = \alpha q_n. \quad [1]$$

Note that the error probability for a given item depends on the number of times it has been reduced by the factor α ; that is, the number of times it has been presented. Learning is the gradual reduction in the probability of error by repeated presentations of items. This model is sometimes called the *linear model* because the equation describing change in response probability is linear.

In the second model, mastery of an item is not at all gradual. At any point in time a student is in one of two states with respect to each item: the learned state or the unlearned state. If an item in the learned state is presented, the correct response is always given; if an item is in the unlearned state, an incorrect response is given unless the student makes a correct response by guessing. When an unlearned item is presented, it may move into the learned state with probability c . Stated as a difference equation,

$$q_{n+1} = \begin{cases} q_n, & \text{with probability } 1 - c \\ 0, & \text{with probability } c. \end{cases} \quad [2]$$

Once an item is learned, it remains in the learned state throughout the course of instruction. Some items are learned the first time they are presented; others may be presented several times before they are finally learned. Therefore, the list as a whole is learned gradually. But for any particular item, the transition from the unlearned to the learned state occurs on a single trial. The model is sometimes called the *all-or-none model* because of this characterization of the possible states of learning.

The third model to be considered is the *random-trial increments model* and represents a compromise between the linear and all-or-none model (Norman, 1964). For this model:

$$q_{n+1} = \begin{cases} q_n, & \text{with probability } 1 - c \\ \alpha q_n, & \text{with probability } c. \end{cases} \quad [3]$$

If $c = 1$, the random-trial increments model reduces to the linear model; if $\alpha = 0$, it reduces to the all-or-none model. However, for $c < 1$ and $\alpha > 0$, the random-trial increments model generates predictions that are quite distinct from both the linear and the all-or-none models.

For all three models the probability of an error on the first trial is a parameter that may need to be estimated in certain situations; to emphasize this point the initial error probability is written as q' henceforth. It should be noted that the all-or-none model and the random-trial increments model are *response sensitive* in that the learner's particular history of correct and incorrect responses makes a difference in predicting performance on the next presentation of an item. In contrast, the linear model is *response insensitive*; its prediction depends only on the number of prior presentations and is not improved by a knowledge of the learner's response history.

Cost/Benefit Structure

At the present level of analysis, it will expedite matters if some assumptions are made to simplify the appraisal of costs and benefits associated with various strategies. It is tacitly assumed that the subject matter being taught is sufficiently beneficial to justify allocating a fixed amount of time to it for instruction. Since the exercise formats and the time allocated to instruction are the same for all strategies, it is reasonable to assume that the costs of instruction are the same for all strategies as well. If the costs of instruction are equal for all strategies, then for purposes of comparison they may be ignored and attention focused on the comparative benefits of the various strategies. This is an important simplification because it affects the degree of precision necessary in the assessment of costs and benefits. If both costs and benefits are significantly variable in a problem, then it is essential that both quantities 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. This is usually fairly easy to accomplish. In the present problem, for example, it is reasonable to consider all the items equally beneficial. This implies that benefits depend only on the overall probability of a correct response, not on the particular items known. It turns out that this specification of cost and benefit is sufficient for the learning models to determine optimal strategies.

The above cost/benefit assumptions permit us to concentrate on the main concern of this

article, the derivation of the educational implications of learning models. Also, they are approximately valid in many instructional contexts. Nevertheless, it must be recognized that in the majority of cases these assumptions will not be satisfied. For instance, the assumption that the alternative strategies cost the same to implement usually does not hold. It only holds as a first approximation in the case being considered here. In the present formulation of the problem, a fixed amount of time is allocated for study and the problem is to maximize learning, subject to this time constraint. An alternative formulation that is more appropriate in some situations fixes a minimum criterion level for learning. In this formulation, the problem is to find a strategy for achieving this criterion level of performance in the shortest time. As a rule, both costs and benefits must be weighed in the analysis, and frequently subtopics within a curriculum vary significantly in their importance. Sometimes there is a choice among several exercise formats. In certain cases, whether or not a certain topic should be taught at all is the critical question. Smallwood (1971) has treated a problem similar to the one considered in this article in a way that includes some of these factors in the structure of costs and benefits.

Deducing Strategies from the Learning Models

Optimal strategies can be deduced for the linear and all-or-none models under the assumption that all items have the same learning parameters and initial error probabilities. The situation is more complicated in the case of the random-trial increments model. An approximation to the optimal strategy for the random-trial increments case will be discussed later; in this case the strategy explicitly allows for individual differences in parameter values.

For the linear model, if an item has been presented n times, the probability of an error on the next presentation of the item is $\alpha^{n-1}q'$; when the item is presented, the error probability is reduced to $\alpha^n q'$. The size of the reduction is thus $\alpha^{n-1}(1-\alpha)q'$. Observe that the size of the decrement in error probability gets smaller with each presentation of the item. This observation can be used to deduce that

the following procedure is optimal:

On a given day, form the sublist of M items by selecting those items that have received the fewest presentations up to that point. If more than M items satisfy this criterion, then select items at random from the set satisfying the criterion.

Upon examination, this strategy is seen to be equivalent to the standard cyclic presentation procedure commonly employed in experiments on paired-associate learning. It amounts to presenting all items once, randomly reordering them, presenting them again, and repeating the procedure until the number of days allocated to instruction have been exhausted.

According to the all-or-none model, once an item has been learned there is no further reason to present it. Since all unlearned items are equally likely to be learned if presented, it is intuitively reasonable that the optimal presentation strategy selects the item least likely to be in the learned state for presentation. In order to discover a good index of the likelihood of being in the learned state, consider a student's response protocol for a single item. If the last response was incorrect, the item was certainly in the unlearned state at that time, although it may then have been learned during the study period that immediately followed. If the last response was correct, then it is more likely that the item is now in the learned state. In general, the more correct responses there are in the protocol since the last error on the item, the more likely it is that the item is in the learned state.

The preceding observations provide a heuristic justification for an algorithm which Karush and Dear (1966) have proved is in fact the optimal strategy for the all-or-none model. The optimal strategy requires that for each student a bank of counters be set up, one for each word in the list. To start, M different items are presented each day until each item has been presented once and a 0 has been entered in its counter. On all subsequent days, the strategy requires that we conform to the following two rules:

1. Whenever an item is presented, increase its counter by 1 if the subject's response is correct, but reset it to 0 if the response is incorrect.
2. Present the M items whose counters are lowest among all items. If more than M items are eligible, then select randomly as many items as are needed to complete the sublist of size M from those having the same

highest counter reading, having selected all items with lower counter values.

For example, suppose six items are presented each day, and after a given day a certain student has four items whose counters are 0, four whose counters are 1, and higher values for the rest of the counters. His study list would consist of the four items whose counters are 0, and two items selected at random from the four whose counters are 1.

It has been possible to find relatively simple optimal strategies for the linear and all-or-none models. It is noteworthy that neither strategy depends on the values of the parameters of the respective models (i.e., on α , c , or q'). Another exceptional feature of these two models is that it is possible to condense a student's response protocol to one index per item without losing any information relevant to presentation decisions. Such condensations of response protocols are referred to as *sufficient histories* (Groen & Atkinson, 1966). Roughly speaking, an index summarizing the information in a student's response protocol is a sufficient history if any additional information from the protocol would be redundant in the determination of the student's state of learning. The concept is analogous to a *sufficient statistic*. If one takes a sample of observations from a population with an underlying normal distribution and wishes to estimate the population mean, the sample mean is a sufficient statistic. Other statistics that can be calculated (such as the median, the range, and the standard deviation) cannot be used to improve on the sample mean as an estimate of the population mean, though they may be useful in assessing the precision of the estimate. In statistics, whether or not data can be summarized by a few simple sufficient statistics is determined by the nature of the underlying distribution. For educational applications, whether or not a given instructional process can be adequately monitored by a simple sufficient history is determined by the model representing the underlying learning process.

The random-trial increments model appears to be an example of a process for which the information in the subject's response protocol cannot be condensed into a simple sufficient history. It is also a model for which the optimal strategy depends on the values of the model

parameters. Consequently, it is not possible to state a simple algorithm for the optimal presentation strategy for this model. Suffice it to say that there is an easily computable formula for determining which item has the best expected immediate gain, if presented. The strategy that presents this item should be a reasonable approximation to the optimal strategy (Calfee, 1970). More is said later regarding the problem of parameter estimation and some of its ramifications.

If the three models under consideration are to be ranked on the basis of their ability to account for data from laboratory experiments employing the standard presentation procedure, the order of preference is clear. The all-or-none model provides a better account of the data than the linear model, and the random-trial increments model is better than either of them (Atkinson & Crothers, 1964). This does not necessarily imply, however, that the optimization strategies derived from these models will receive the same ranking. The standard cyclic presentation procedure used in most learning experiments may mask certain deficiencies in the all-or-none or random-trial increments models which would manifest themselves when the optimal presentation strategy specified by one or the other of these models was employed.²

² This type of result was obtained by Dear, Silberman, Estavan, and Atkinson (1967). They used the all-or-none model to generate optimal presentation schedules where there were no constraints on the number of times a given item could be presented for test and study within an instructional period. Under these conditions the model generates an optimal strategy that has a high probability of repeating the same item over and over again until a string of correct responses occurs. In their experiment the all-or-none strategy proved quite unsatisfactory when compared with the standard presentation schedule. The problem was that the all-or-none model provides an accurate account of learning when the items are well spaced, but fails badly under highly massed conditions. Laboratory experiments prior to the Dear et al. (1967) study had not employed a massing procedure, and this particular deficiency of the all-or-none model had not been made apparent. The important remark here is that the analysis of instructional problems can provide important information in the development of learning models. In certain cases the set of laboratory tasks that the psychologist deals with may be such that it fails to uncover that particular condition which would cause the model to fail. By analyzing optimal learning conditions we are imposing a somewhat different test on a learning model, which may provide a more sensitive measure of its adequacy.

EVALUATION OF THE ALL-OR-NONE STRATEGY

Lorton³ compared the all-or-none strategy with the standard procedure in an experiment in computer-assisted spelling instruction with elementary school children. The former strategy is optimal if the learning process is indeed all-or-none, whereas the latter is optimal if the process is linear. The experiment was one phase of the Stanford Reading Project using computer facilities at Stanford University linked via telephone lines to student terminals in the schools.

Individual lists of 48 words were compiled in an extensive pretest program to guarantee that each student would be studying words of approximately equal difficulty, which he did not already know how to spell. A within-subjects design was used in an effort to make the comparison of strategies as sensitive as possible. Each student's individualized list of 48 words was used to form two comparable lists of 24 words, one to be taught using the all-or-none strategy and the other using the standard procedure.

Each day a student was given training on 16 words, 8 from the list for standard presentation and 8 from the list for presentation according to the all-or-none strategy. There were 24 training sessions followed by 3 days for testing all the words; approximately 2 weeks later 3 more days were spent on a delayed retention test. Using this procedure, all words in the standard presentation list received exactly one presentation in successive 3-day blocks during training. Words in the list presented according to the all-or-none algorithm received from zero to three presentations in successive 3-day blocks during training, with one presentation being the average. A flow chart of the daily routine is given in Figure 1.

The results of the experiment are summarized in Figure 2. The proportions of correct responses are plotted for successive 3-day blocks during training, followed by the first overall test and then the 2-week delayed test. Note that during training the proportion correct is always *lower* for the all-or-none procedure than for the standard procedure, but on

³ P. Lorton, Jr. Computer-based instruction in spelling: An investigation of optimal strategies for presenting instructional material. Unpublished doctoral dissertation, Stanford University, 1972.

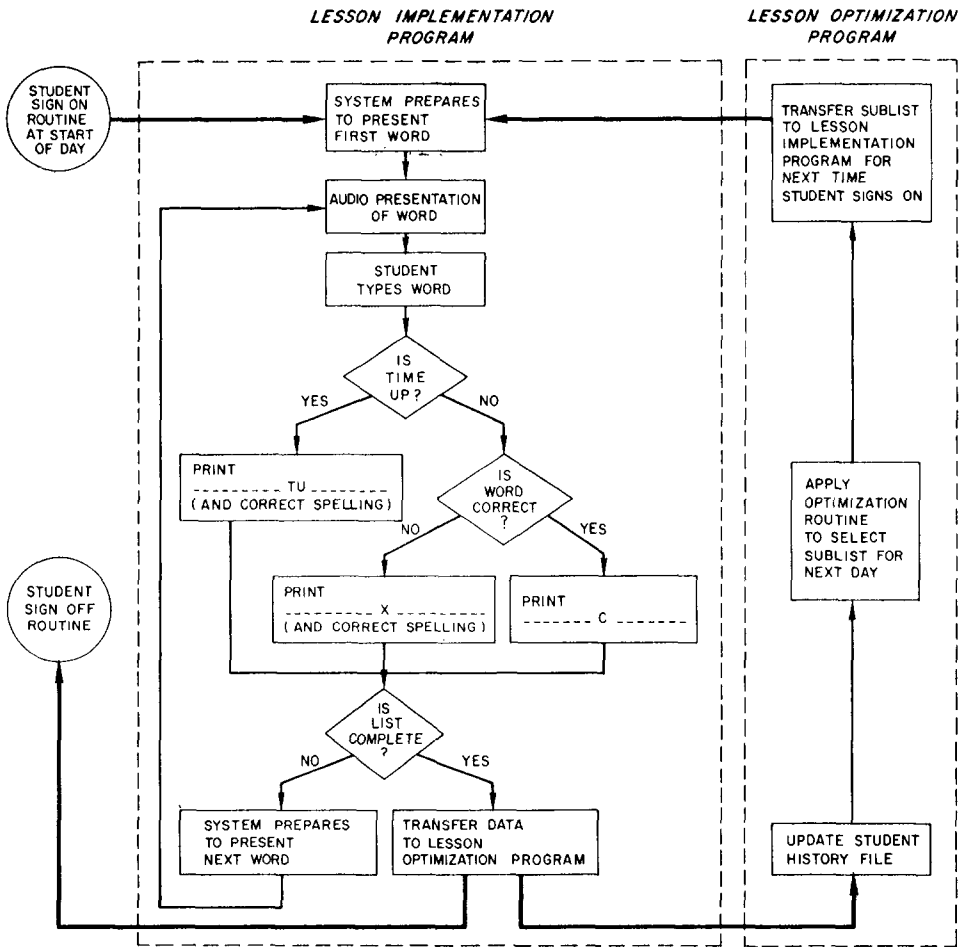


FIG. 1. Daily list presentation routine.

both the final test and the retention test the proportion correct is *greater* for the all-or-none strategy. Analysis of variance tests verified that these results are statistically significant. The advantage of approximately 10 percentage points on the posttests for the all-or-none procedure is of practical significance as well.

The observed pattern of results is exactly what would be predicted if the all-or-none model does indeed describe the learning process. As was shown earlier, final test performance should be better when the all-or-none optimization strategy is adopted as opposed to the standard procedure. Also the greater proportion of error for this strategy during training is to be expected. The all-or-none strategy presents the items least likely to be in

the learned state, so it is natural that more errors would be made during training.

TEST OF A PARAMETER-DEPENDENT STRATEGY

As noted earlier, the strategy derived for the all-or-none model in the case of homogeneous items does not depend on the actual values of the model parameters. In many situations either the assumptions of the all-or-none model or the assumption of homogeneous items or both are seriously violated, so it is necessary to consider strategies based on more general models. Laubsch⁴ considered the optimization

⁴ J. H. Laubsch. An adaptive teaching system for optimal item allocation. Unpublished doctoral dissertation, Stanford University, 1969.

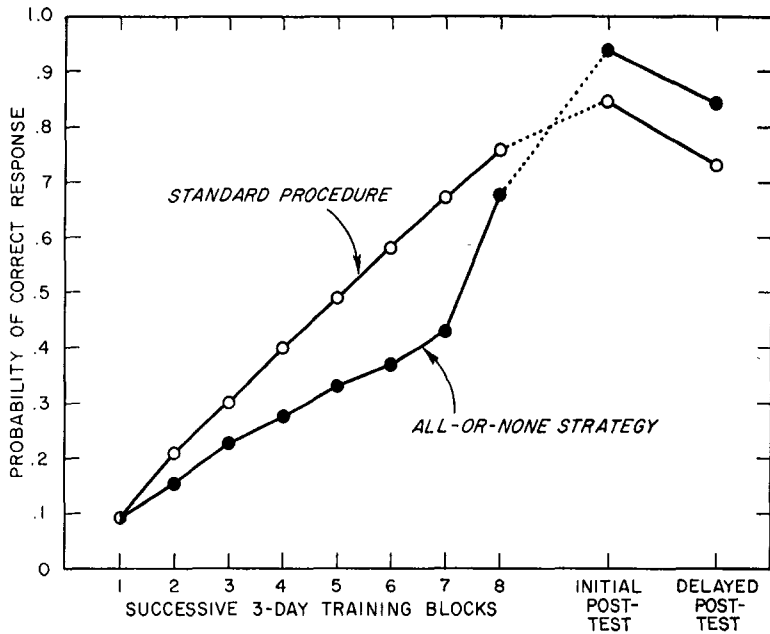


Fig. 2. Probability of correct response in Lorton's experiment.

problem for cases where the random-trial increments model is appropriate. He made what is perhaps a more significant departure from the assumptions of the all-or-none strategy by allowing the parameters of the model to vary with students and items. The following discussion is based upon Laubsch's work, but introduces a more satisfactory formulation of individual differences. This change and the estimation of initial condition parameters produce experimental measures of the effectiveness of optimization procedures that are significantly greater than those reported by Laubsch.

It is not difficult to derive an approximation to the optimal strategy for the random-trial increments model that can accommodate student and item differences in parameter values, if these parameters are known. Since parameter values must be specified in order to make the necessary calculations to determine the optimal study list, it makes little difference whether these numbers are fixed or vary with students and items. However, making estimates of these parameter values in the heterogeneous case presents some difficulties.

When the parameters of a model are homogeneous, it is possible to pool data from

different subjects and items to obtain precise estimates. Estimates based on a sample of students and items can be used to predict the performance of other students or the same students on other items. When the parameters are heterogeneous, these advantages no longer exist unless variations in the parameter values take some known form. For this reason it is necessary to formulate a model stating the composition of each parameter in terms of a subject and item component.

Let π_{ij} be a generic symbol for a parameter characterizing student i and item j . An example of the kind of relationship desired is a fixed-effects Subjects \times Items analysis of variance model:

$$E(\pi_{ij}) = m + a_i + d_j \quad [4]$$

where m is the mean, a_i is the ability of student i , and d_j is the difficulty of item j . Because the learning model parameters we are interested in are probabilities, the above assumption of additivity is not met; that is, there is no guarantee that Equation 4 would yield estimates bounded between 0 and 1. But there is a transformation of the parameter that circumvents this difficulty. In the present context,

this transformation has an interesting intuitive justification.

Instead of thinking directly in terms of the parameter π_{ij} , it is helpful to think in terms of the odds ratio $\pi_{ij}/1 - \pi_{ij}$. Allow two assumptions: (a) the odds ratio is proportional to student ability; (b) the odds ratio is inversely proportional to item difficulty. This can be expressed algebraically as

$$\frac{\pi_{ij}}{1 - \pi_{ij}} = \kappa \frac{a_i}{d_j}, \quad [5]$$

where κ is a proportionality constant. Taking logarithms on both sides yields

$$\log \frac{\pi_{ij}}{1 - \pi_{ij}} = \log \kappa + \log a_i - \log d_j. \quad [6]$$

The logarithm of the odds ratio is usually referred to as the "logit." Let $\log \kappa = \mu$, $\log a_i = A_i$, and $-\log d_j = D_j$. Then Equation 6 becomes

$$\text{logit } \pi_{ij} = \mu + A_i + D_j. \quad [7]$$

Thus, the two assumptions made above lead to an additive model for the values of the parameters transformed by the logit function. Equation 7, by defining a subject-item parameter π_{ij} in terms of a subject parameter A_i applying to all items and an item parameter D_j applying to all subjects, significantly reduces the number of parameters to be estimated. If there are N items and S subjects, then the model requires only $N + S$ parameters to specify the learning parameters for $N \times S$ subject items. More importantly, it makes it possible to predict a student's performance on items he has not been exposed to from the performance of other students on them. This formulation of learning parameters is essentially the same as the treatment of an analogous problem in item analysis given by Rasch (1966). Discussion of this and related models for problems in mental test theory is given by Birnbaum (1968).

Given data from an experiment, Equation 7 can be used to obtain reasonable parameter estimates, even though the parameters vary with students and items. The parameters π_{ij} are first estimated for each student-item protocol, yielding a set of initial estimates.

Next the logistic transformation is applied to these initial estimates, and then using these values subject and item effects (A_i and D_j) are estimated by standard analysis of variance procedures. The estimates of student and item effects are used to adjust the estimate of each transformed student-item parameter, which in turn is transformed back to obtain the final estimate of the original student-item parameter.

The first students in an instructional program that employs a parameter-dependent optimization scheme like the one outlined above do not benefit maximally from the program's sensitivity to individual differences in students and items; the reason is that the initial parameter estimates must be based on the data from these students. As more and more students complete the program, estimates of the D_j s becomes more precise until finally they may be regarded as known constants of the system. When this point has been reached, the only task remaining is to estimate A_i for each new student entering the program. Since the D_j s are known, the estimates of π_{ij} for a new student are of the right order, although they may be systematically high or low until the student component can be accurately assessed.

Parameter-dependent optimization programs with the adaptive character just described are potentially of great importance in long-term instructional programs. Of interest here is the random-trial increments model, but the method of decomposing parameters into student and item components would apply to other models as well. We turn now to an experimental test of the adaptive optimization program based on the random-trial increments model. In this case the parameters α , c , and q' of the random-trial increments model were separated into item and subject components following the logic of Equation 7. That is, the parameters for subject i working on item j were defined as follows:

$$\begin{aligned} \text{logit } \alpha_{ij} &= \mu^{(\alpha)} + A_i^{(\alpha)} + D_j^{(\alpha)} \\ \text{logit } c_{ij} &= \mu^{(c)} + A_i^{(c)} + D_j^{(c)} \\ \text{logit } q'_{ij} &= \mu^{(q')} + A_i^{(q')} + D_j^{(q')}. \end{aligned} \quad [8]$$

Note that $A_i^{(\alpha)}$, $A_i^{(c)}$, and $A_i^{(q')}$ are measures of the ability of subject i and hold for all

items, whereas $D_j^{(a)}$, $D_j^{(c)}$, and $D_j^{(a')}$ are measures of the difficulty of item j and hold for all subjects.

The instructional program was designed to teach 300 Swahili vocabulary items to college-level students. Two presentation strategies were employed: (a) the all-or-none procedure and (b) the adaptive optimization procedure based on the random-trial increments model. As in the Lorton study, a within-subjects design was employed in order to provide a sensitive comparison of the strategies. For each student two sublists of 150 items were formed at random from the master list; instruction on items from one sublist was governed by the all-or-none strategy and by the adaptive optimization strategy for the other sublist. Each day a student was tested on and studied 100 items presented in a random order; 50 items were from the all-or-none sublist chosen using the all-or-none strategy, and 50 from the adaptive optimization list selected according to that strategy. A Swahili word would be displayed, and the student was required to give its English translation. Reinforcement consisted of a printout of the correct Swahili-English pair. Twenty such training sessions were involved, each lasting for approximately 1 hour. Two or three days after the last training session an initial posttest was administered over all 300 items; a delayed posttest was given approximately 2 weeks later.

The lesson optimization program for the random-trial increments model was more complex than those described earlier. Each night the response data for that day were entered into the system and used to update parameter estimates; in this case an exact record of the complete presentation sequence and response history had to be preserved. A computer-based search algorithm was used to estimate parameters and thus the more accurate the previous day's estimates, the more rapid was the search for the updated parameter values. Once updated estimates had been obtained, they were entered into the optimization program to select individual lists for each student to be run the next day. Early in the experiment (before estimates of $D^{(a)}$, $D^{(c)}$, and $D^{(a')}$ had stabilized) the computation time was fairly lengthy, but it rapidly decreased as more data accumu-




FIG. 3. Posttest performance for the all-or-none strategy and for the parameter-dependent strategy of the random-trial increments model.

lated, and the system homed in on precise estimates of item difficulty.

Figure 3 presents the final test results and indicates that for both the initial and delayed posttests the parameter-dependent strategy of the random-trial increments model was markedly superior to the all-or-none strategy; on the initial posttest the relative improvement was 41% and 67% on the delayed posttest. It is apparent that the parameter-dependent strategy was more sensitive than the all-or-none strategy in identifying and presenting those items that would benefit most from additional study. Another feature of the experiment was that students were run in successive groups, each starting about one week after the prior group. As the theory would predict, the overall gains produced by the parameter-dependent strategy increased from one group to the next. The reason is that early in the experiment estimates of item difficulty were crude, but improved with each successive wave

of students. Near the end of the experiment estimates of item difficulty were quite exact, and the only task that remained when a new student came on the system was to estimate his particular $A^{(\alpha)}$, $A^{(c)}$, and $A^{(q')}$ values.

CONCLUDING REMARKS

The studies reported here illustrate one approach that can contribute to the development of a theory of instruction. This is not to suggest that the strategies they tested represent a complete solution to the problem of optimal item selection. The models upon which these strategies are based ignore several potentially important factors, such as short-term memory effects, interitem relationships, and motivation. Undoubtedly, strategies based on learning models that take some of these variables into account would be superior to those analyzed so far.

The task and learning models considered in this article are extremely simple and of restricted generality; nevertheless, there are at least two reasons for studying them. First, this type of task occurs in many fields of instruction and needs to be understood in its own right. No matter what the pedagogical orientation, it is hard to conceive of an initial reading program or foreign language course that does not involve some form of list-learning activity. In this respect it should be noted that the parameter-dependent strategy of the random-trial increments model has been incorporated into several parts of the Stanford computer-assisted instruction program in initial reading; although no formal evaluation has yet been made, the strategy appears to be highly effective.

There is a second reason for the type of analysis reported here. By making a careful study of a few cases that can be understood in detail, it is possible to develop prototypical procedures for analyzing more complex optimization problems. At present, analyses comparable to those reported here cannot be made for many problems of central interest to education, but by having examples of the above sort it is possible to specify with more clarity the steps involved in devising optimal procedures. Three aspects need to be emphasized: (a) the development of an adequate description of the

learning process, (b) the assessment of costs and benefits associated with possible instructional actions and states of learning, and (c) the derivation of optimal strategies based on the goals set for the student. The examples considered here deal with each of these factors and point out the issues that arise.

It has become fashionable in recent years to criticize learning theorists for ignoring the prescriptive aspects of instruction, and some have argued that efforts devoted to the laboratory analysis of learning should be redirected to the study of learning as it occurs in real-life situations. These criticisms are not entirely unjustified for in practice psychologists have too narrowly defined the field of learning, but to focus all effort on the study of complex instructional tasks would be a mistake. Some successes might be achieved, but, in the long run, understanding complex learning situations must depend upon a detailed analysis of the elementary perceptual and cognitive processes that make up the human information handling system. The trend to press for relevance of learning theory is healthy, but if the surge in this direction goes too far, we will end up with a massive set of prescriptive rules and no theory to integrate them.

It needs to be emphasized that the interpretation of complex phenomena is problematical, even in the best of circumstances. The case of hydrodynamics is a good example for it is one of the most highly developed branches of theoretical physics. Differential equations expressing certain basic hydrodynamic relationships were formulated by Euler in the eighteenth century. Special cases of these equations sufficed to account for a wide variety of experimental data. These successes prompted Lagrange to assert that the success would be universal were it not for the difficulty in integrating Euler's equations in particular cases. Lagrange's view is still widely held, in spite of numerous experiments yielding anomalous results. Euler's equations have been integrated in many cases, and the results were found to disagree dramatically with observation, thus contradicting Lagrange's assertion. The problems involve more than mere fine points, and raise serious paradoxes when extrapolations are made from results obtained in the laboratory to actual conditions.

The following quotation from Birkhoff (1960) should strike a sympathetic chord among those trying to relate psychology and education:

These paradoxes have been the subject of many witticisms. Thus, it has been said that in the nineteenth century, fluid dynamicists were divided into hydraulic engineers who observed what could not be explained, and mathematicians who explained things that could not be observed. It is my impression that many survivors of both species are still with us [p. 26].

Research on learning appears to be in a similar state. Educational researchers are concerned with experiments that cannot be readily interpreted in terms of learning theoretic concepts, while psychologists continue to develop theories that seem to be applicable only to the phenomena of the laboratory. Hopefully, work of the sort described here will bridge this gap and help lay the foundations for a theory of instruction.

REFERENCES

- ATKINSON, R. C. Computer-assisted learning in action. *Proceedings of the National Academy of Sciences*, 1969, **63**, 588-594.
- ATKINSON, R. C., & CROTHERS, E. J. A comparison of paired-associate learning models having different acquisition and retention axioms. *Journal of Mathematical Psychology*, 1964, **2**, 285-315.
- ATKINSON, R. C. & FLETCHER, J. D. Teaching children to read with a computer. *The Reading Teacher*, 1972, **25**, 319-327.
- BELLMAN, R. *Adaptive control processes*. Princeton: Princeton University Press, 1961.
- BIRKHOFF, G. *Hydrodynamics: A study in logic, fact, and similitude*. (2nd ed.) Princeton: Princeton University Press, 1960.
- BIRNBAUM, A. Some latent trait models and their use in inferring an examinee's ability. In F. M. Lord & M. R. Novick (Eds.), *Statistical theories of mental test scores*. Reading, Mass.: Addison-Wesley, 1968.
- BRUNER, J. S. Some theorems on instruction stated with reference to mathematics. In E. R. Hilgard (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.
- CALFEE, R. C. The role of mathematical models in optimizing instruction. *Scientia: Revue Internationale de Synthèse Scientifique*, 1970, **105**, 1-25.
- 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.
- GROEN, G. J., & ATKINSON, R. C. Models for optimizing the learning process. *Psychological Bulletin*, 1966, **66**, 309-320.
- 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.
- KARUSH, W., & DEAR, R. E. Optimal stimulus presentation strategy for a stimulus sampling model of learning. *Journal of Mathematical Psychology*, 1966, **3**, 19-47.
- NORMAN, M. F. Incremental learning on random trials. *Journal of Mathematical Psychology*, 1964, **2**, 336-350.
- RAIFFA, H., & SCHLAIFFER, R. *Applied statistical decision theory*. Cambridge, Mass.: M.I.T. Press, 1968.
- RASCH, G. An individualistic approach to item analysis. In P. F. Lazarsfeld & N. W. Henry (Eds.), *Readings in mathematical social science*. Chicago: Science Research Associates, 1966.
- SMALLWOOD, R. D. The analysis of economic teaching strategies for a simple learning model. *Journal of Mathematical Psychology*, 1971, **8**, 285-301.

(Received October 15, 1970)