

AN APPROACH TO THE PSYCHOLOGY OF INSTRUCTION

by

R. C. Atkinson and J. A. Paulson

TECHNICAL REPORT NO. 157

August 14, 1970

PSYCHOLOGY SERIES

Reproduction in Whole or in Part is Permitted for
any Purpose of the United States Government

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

UNIVERSITY OF CALIFORNIA LIBRARY
OF

BERKELEY, CALIFORNIA

1968

1968

1968

UNIVERSITY OF CALIFORNIA LIBRARY

BERKELEY, CALIFORNIA

UNIVERSITY OF CALIFORNIA LIBRARY

BERKELEY, CALIFORNIA

1968

AN APPROACH TO THE PSYCHOLOGY OF INSTRUCTION¹

R. C. Atkinson and J. A. Paulson

Stanford University
Stanford, California 94305

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 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 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. Recent 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 paper will examine 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 which 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 which must be specified are the following:

- 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 doesn't 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 paper. 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 paper 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 shall discuss, the successful strategy produced an unusually 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 WHICH ARISES IN COMPUTER-ASSISTED INSTRUCTION

One application of computer-assisted instruction (CAI) which 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). The curriculum materials in such programs frequently take the form of lists of instructional units or items. The objective of the CAI 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, Chetin, & Stauffer, 1970). The vocabularies of several of the commonly used basal readers were compiled into one dictionary and a variety of exercises using these words was designed to develop reading skills. Separate exercise formats were designed 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 format to another, but most include a sequence in which an 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. If the student types the correct word, he receives a reinforcing message and proceeds to the next presentation.

If he 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 one version of the program, items are 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 recent dissertation studies to be described below are concerned with the development of such strategies. Lorton (1969) studied alternative presentation strategies for teaching spelling words in an experiment with elementary school children, and Laubsch (1969) studied similar strategies for teaching Swahili vocabulary items to Stanford undergraduates.

The optimization problems in both the Lorton and Laubsch 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 item is presented only once for test followed immediately by a brief study period. The total set of N items is extremely large with regard to the 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 the posttest by an appropriate selection of sublists from day to day. The strategy for selecting sublists is

dynamic (or response sensitive, using the terminology of Groen and Atkinson, 1966) to the extent that it depends upon the student's history of performance.

Three Models of the Learning Process

Two extremely simple learning models will be considered first. Then a third model which combines features of the first two will be 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 presented. In terms of the classification scheme introduced by Groen and Atkinson (1966), the process is response-insensitive. The state of the learner is related to his responses as follows: at the start of the experiment, all items have some initial probability of error, say q_1 ; each time an item is presented, its error probability is reduced by a factor α , which is less than one. Stated as an equation, this becomes

$$(1) \quad q_{n+1} = \alpha q_n,$$

or alternatively,

$$(2) \quad q_{n+1} = \alpha^n q_1.$$

The error probability for a given item depends on the number of times it has been reduced by the factor α ; i.e., 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 (Bush & Mosteller, 1955).

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 an equation,

$$(3) \quad q_{n+1} = \begin{cases} q_n, & \text{with probability } 1-c \\ 0, & \text{with probability } c. \end{cases}$$

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 (Atkinson & Crothers, 1964).

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

$$(4) \quad q_{n+1} = \begin{cases} q_n, & \text{with probability } 1-c \\ \alpha q_n, & \text{with probability } c. \end{cases}$$

If $c = 1$, then $q_{n+1} = \alpha q_n$ and the model reduces to the linear model. If $\alpha = 0$, then the model reduces to the all-or-none model. However, if $c < 1$ and $\alpha > 0$, the RTI model generates predictions that are quite distinct from both the linear and the all-or-none models. It should be noted that both the all-or-none model and the RTI model are response sensitive in the sense that the learner's particular history of correct and incorrect responses makes a difference in predicting performance on the next presentation of an item.

The 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 important 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 vocabulary items equally important. 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 models to determine optimal strategies.

The above cost/benefit assumptions permit us to concentrate on the main concern of this paper, 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 which 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 (1970) has treated a problem similar to the one considered in

this paper 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. The situation is more complicated in the case of the RTI model. An approximation to the optimal strategy for the RTI case will be discussed; the strategy will explicitly allow for 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 $\alpha^{n-1}q_1$; when the item is presented, the error probability is reduced to $\alpha^n q_1$. The size of the reduction is thus $\alpha^{n-1}(1-\alpha)q_1$. 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 verbal 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 6 items are presented each day and after a given day a certain student has 4 items whose counters are 0, 4 whose counters are 1, and higher values for the rest of the counters. His study list would consist of the 4 items whose counters are 0, and 2 items selected at random from the 4 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_1). 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. More will be 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 RII models which would manifest themselves when the optimal presentation strategy specified by one or the other of these models was employed.²

AN EVALUATION OF THE ALL-OR-NONE STRATEGY

Lorton (1969) 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 three days for testing all the words; approximately two weeks later three 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 0 to 3 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. Special features of the lesson implementation program allowed students to correct typing errors or request repetition of audio messages before a response was evaluated. These features reduced the likelihood of missing a word because of momentary inattention or typing errors.

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 two-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 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 ten 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

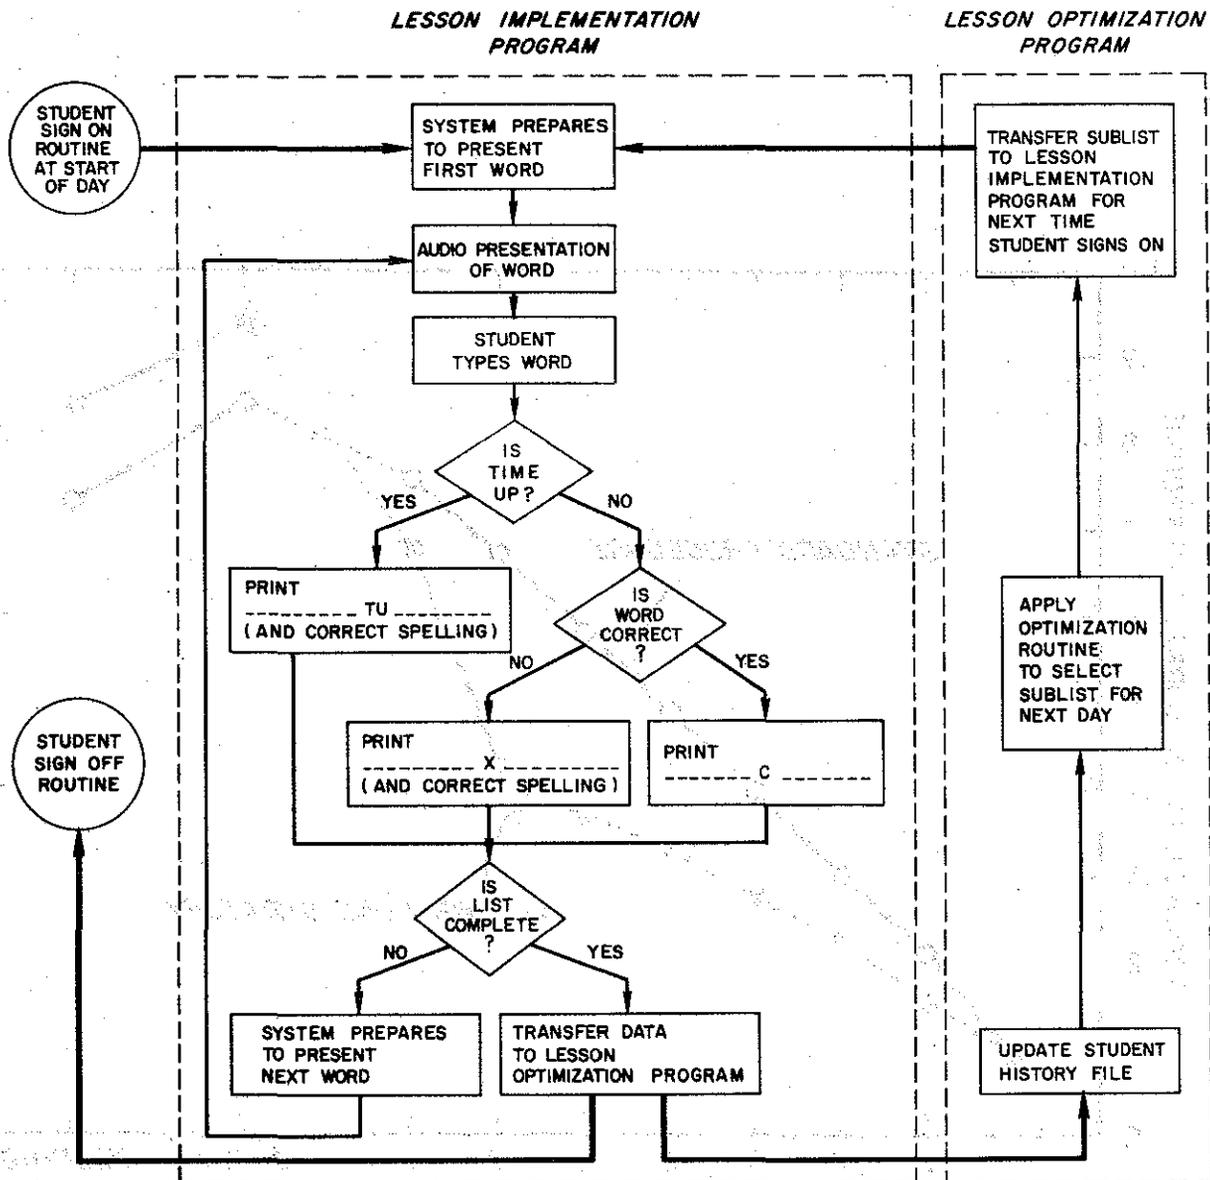


Figure 1. Daily list presentation routine.

WOLFEHART 1964

WOLFEHART 1964

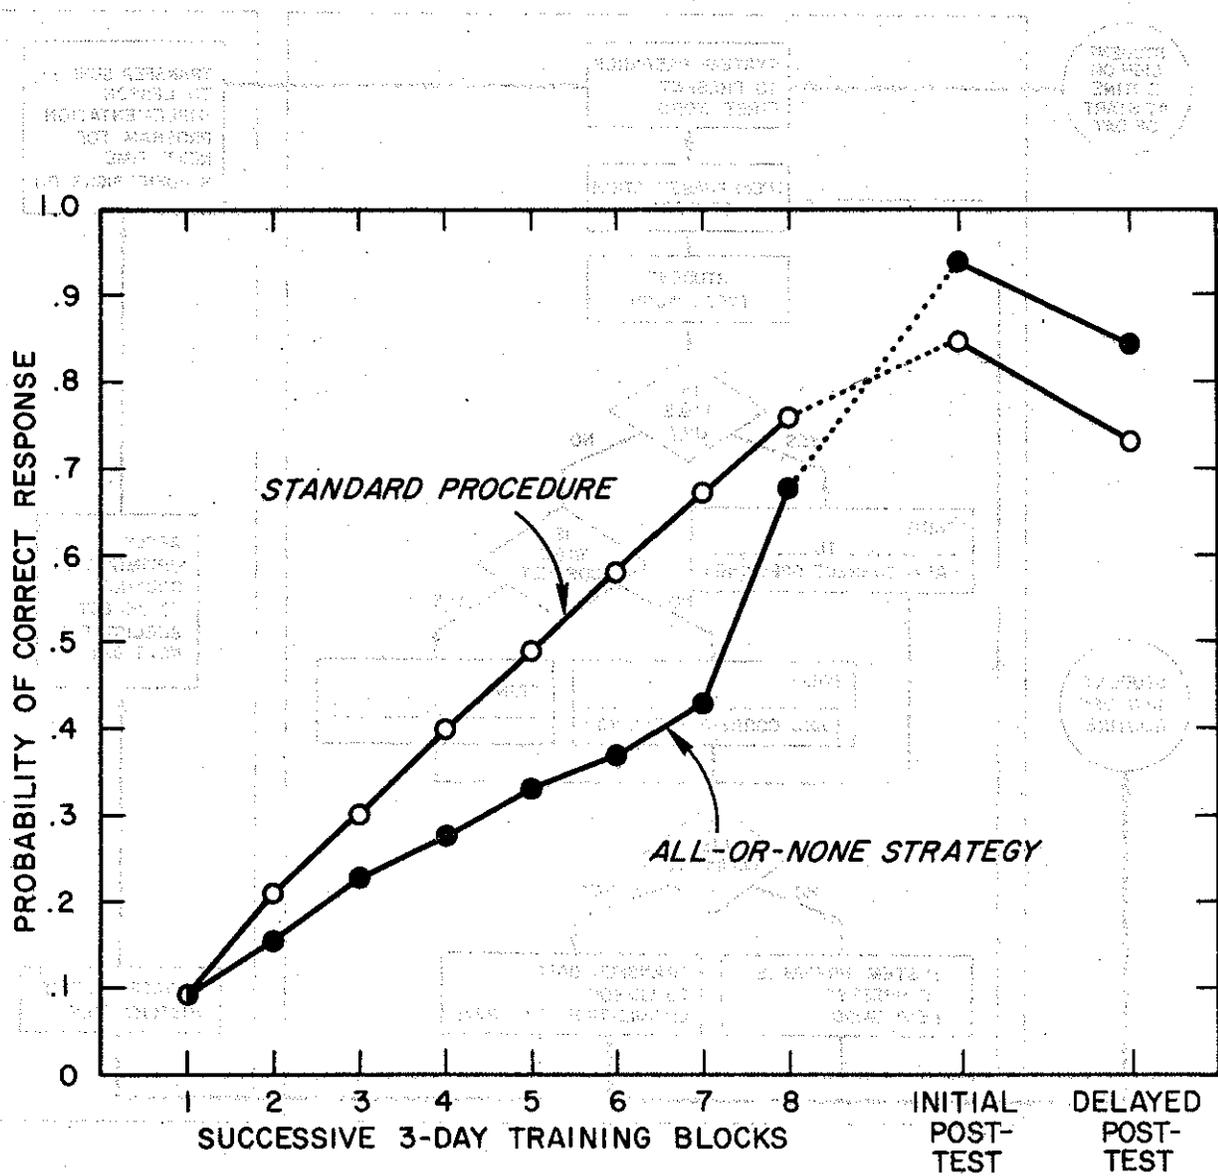


Figure 2. Probability of correct response in Lorton's experiment.

least likely to be in the learned state, so it is natural that more errors would be made during training. Thus, according to the all-or-none model the most rapid learning results from a routine which, in a sense, maximizes the student's failures during training. This apparent anomaly will be considered later.

A 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 other models. Laubsch (1969) considered the optimization problem for cases where the RTI 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.

It is not difficult to derive an approximation to the optimal strategy for the RTI 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. The model suggested here is a simplification of the procedure Laubsch employed.

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-by-items analysis of variance model:

$$(5) \quad E(\pi_{ij}) = m + a_i + d_j$$

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 Eq. 5 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: (1) the odds ratio is proportional to student ability; (2) the odds ratio is inversely proportional to item difficulty. This can be expressed algebraically as

$$(6) \quad \frac{\pi_{ij}}{1-\pi_{ij}} = \kappa \frac{a_i}{d_j}$$

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

$$(7) \quad \log \frac{\pi_{ij}}{1-\pi_{ij}} = \log K + \log a_i - \log d_j$$

The logarithm of the odds ratio is usually referred to as the "logit."

Let $\log K = \mu$, $\log a_i = A_i$, and $-\log d_j = D_j$. Then Eq. 7 becomes

$$(8) \quad \text{logit } \pi_{ij} = \mu + A_i + D_j$$

Thus, the two assumptions made above lead to an additive model for the values of the parameters transformed by the logit function. Equation 8, 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, Eq. 8 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 which 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 become 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 RTI model, but the method of decomposing parameters into student and item components would apply to other models as well. We turn now to Laubsch's experimental test of the adaptive optimization program based on the RTI model. In this case both parameters α and c of the RTI model were separated into item and subject components following the logic of Eq. 8. That is, the parameters for subject i working on item j were defined as follows:

$$(9) \quad \text{logit } \alpha_{ij} = \mu^{(\alpha)} + A_i^{(\alpha)} + D_j^{(\alpha)}$$

$$\text{logit } c_{ij} = \mu^{(c)} + A_i^{(c)} + D_j^{(c)}$$

Note that $A_i^{(\alpha)}$ and $A_i^{(c)}$ are measures of the ability of subject i and hold for all items, whereas $D_j^{(\alpha)}$ and $D_j^{(c)}$ are measures of the difficulty of item j and hold for all subjects.

The instructional program was designed to teach 420 Swahili vocabulary items to undergraduate students at Stanford University. Three presentation strategies were employed: (1) the standard cyclic procedure, (2) the all-or-none procedure, and (3) the adaptive optimization procedure based on the RTI model. As in the Lorton study, a within-subjects design was employed in order to provide a sensitive comparison of the strategies. The procedural details were essentially the same as in Lorton's experiment, except for the fact that 14 training sessions were involved, each lasting for approximately one hour. A Swahili word would be presented and a response set of five English words would appear on the teletype. The student's task was to type the number of the correct alternative. Reinforcement consisted of a "+" or "-" and a printout of the correct Swahili-English pair.

The lesson optimization program for the RTI model was more complex than those described earlier. Each night the response data for that day was entered into the system and used to update estimates of the α 's and c 's; 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 item sublists for each student to be run the next day. Early in the experiment (before estimates of the $D^{(\alpha)}$'s and $D^{(c)}$'s had stabilized) the computation time was fairly lengthy, but it rapidly decreased as more data accumulated and the system homed in on precise estimates of item difficulty.

The results of the experiment favored the parameter-dependent strategy for both a final test administered immediately after the termination of instruction and for a delayed retention test presented several weeks later. Stated otherwise, the parameter-dependent strategy of the RTI model was more sensitive than the all-or-none or linear strategies in identifying and presenting those items that would benefit most from additional training. 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 improve 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 $A^{(\alpha)}$ and $A^{(c)}$ values.

IMPLICATIONS FOR FURTHER RESEARCH

The studies of both Laubsch and Lorton 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, inter-item 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 studies described here avoided many difficulties associated with short-term retention effects by presenting items for test and study at most once per day. But in many situations it is desirable to employ procedures in which items can be presented more than once per day. If such procedures are employed, experiments by Greeno (1964), Fishman, Keller, and Atkinson (1968), and others indicate that the optimal strategy will have to take short-term memory effects into account. The results reported by these investigators can be accounted for by a more general model similar in many respects to the all-or-none and RTI model (Atkinson & Shiffrin, 1968). The difference is that the more general model has two learned states: a long-term memory state and a short-term state. An item in the long-term state remains there for a relatively indefinite period of time, but an item in the short-term state will be forgotten with a probability that depends on the interval between successive presentations. When items receive repeated presentations in short intervals of time, they may be responded to correctly several times in a row because they are in the short-term state. A strategy (like the one based on the all-or-none model) which did not take this possibility into account would regard these items as well learned and tend

not to present them again, when in fact they would have a high probability of being forgotten.

In many situations some of the items to be presented are interrelated in an obvious way; a realistic model of the learning process would have to reflect these organizational factors. It is likely that the difference between the standard procedure and the best possible procedure is very large in these instances so there is considerable reason to study them. Unfortunately, as yet very little work has been done in formulating mathematical models for such interrelationships, but there are several obvious directions to pursue.

The results of an experiment reported by Hartley (1968) illustrate the complexity of empirical relationships in this area. The study involved the Stanford CAI Project in initial reading and was designed to investigate two types of list organization: minimal versus maximal contrast, combined with three sources of cue; the word itself, the word plus a picture, and the word plus a sentence context cue. Hartley was interested in the relative merit of these conditions for the acquisition of an initial sight-word vocabulary. Fries (1962) had advocated the use of minimal contrast lists in reading instruction in order to exploit linguistic regularities. On the other hand, Rothkopf (1958) found that lists composed of dissimilar items were learned more rapidly than those with small or minimal differences. Hartley's experiment indicated that which list organization is best depends on the cue source. When the word itself was the only cue, performance was best on minimal contrast lists. When the word was augmented with a picture cue, there was little difference in performance on the two kinds of list. But in the presence of a context cue, performance was best on the maximal contrast lists.

In the description of Lorton's experiment we mentioned that the all-or-none strategy produced a higher error rate during learning than the standard procedure. If some observations made by Suppes (1967) are correct, this fact suggests that a better strategy could be devised. Suppes argues that in long-term instructional programs it is crucial to balance considerations of frustration due to material that is too difficult against boredom for material that is too easy. He conjectures that there is an optimal error rate, which if deviated from adversely affects learning. This conjecture poses two interesting problems: first, to determine the range and degree to which it is correct; second, to formulate a model of the learning process that takes account of error rates. The resulting optimization scheme would need to estimate the optimum error rate for each student and these estimates in turn would be inputs to the decision-theoretic problem. The view that there is an optimal error rate is held by many psychologists and educators, so information about this question would be of some significance.

The directions for research which have been discussed here point to the need for considerable theoretical and experimental groundwork to serve as a basis for devising instructional strategies. There are fundamental issues in learning theory that need to be explored and intuitively reasonable strategies of instruction to be tried out. It seems likely that new proposals for optimal procedures will involve parameter-dependent strategies. If this is the case, then provision for variations in parameter values due to differences among students and curriculum materials will be an important consideration. The approach described in the discussion of Laubsch's study could well be applicable to these problems.

CONCLUDING REMARKS

This paper has presented examples of the kind of study we believe can contribute to the psychology of instruction, as distinguished from the psychology of learning. Such studies have both descriptive and prescriptive aspects. Each aspect in turn has an empirical and a theoretical component. The examples described involved the derivation of optimal presentation strategies for fairly simple learning models and the comparison of these strategies in CAI experiments. In both studies the optimal strategy produced significantly better results on criterion tests than a standard cyclic procedure. Evaluation of these experiments suggests a number of ways in which the strategies might be improved, and generalized to a broader range of problems.

The task and learning models considered in this paper 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 different fields of instruction and should 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. Although this type of task has frequently been misused in the design of curricula, its use is so widespread that optimal procedures need to be specified.

There is a second and equally important reason for the type of analysis reported here. By making a study of one case that can be pursued 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 list with more clarity the steps involved in devising optimal procedures. Three aspects need to be emphasized: (1) the development of an adequate description of the learning process, (2) the assessment of costs and benefits associated with possible instructional actions and states of learning, and (3) 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 chide learning theory for ignoring the prescriptive aspects of instruction, and some have even argued that efforts devoted to the laboratory analysis of learning should be redirected to the study of complex phenomena as they occur in instructional 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 initial 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 from which the information handling system of each human being is constructed. 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 but no theory to integrate them. Information processing models of memory and thought and the work on psycholinguistics are promising avenues of research on the learning process, and the prospects are good that they will provide useful theoretical ideas for interpreting the complex phenomena of instruction.

It needs to be emphasized, however, that the interpretation of complex phenomena is problematical, even in the best of circumstances. Consider, for example, the case of hydrodynamics, 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 by many, 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 wind tunnels and from models of harbors and rivers to actual conditions. The following quotation from Birkhoff (1960) should strike a sympathetic cord among those trying to relate psychology and education: "These paradoxes have been the subject of many witticisms. Thus, it has recently 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."

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 theory, while psychologists continue to develop theories that seem to be applicable only to the phenomena observed in their laboratories. Hopefully, work of the sort described here will bridge this gap and help lay the foundations for a viable theory of instruction. If the necessary level of interchange between workers in different disciplines can be developed, the prospects for advancing both psychology and education are good.

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., Chetin, H. C., & Stauffer, C. M. Instruction in initial reading under computer control: The Stanford Project. Technical Report 158, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1970.
- Atkinson, R. C., & Shiffrin, R. M. Human memory: A proposed system and its control processes. In K. W. Spence and J. T. Spence (Eds.), The psychology of learning and motivation: Advances in research and theory, Vol. II. New York: Academic Press, 1968.
- Bellman, R. Adaptive control processes. Princeton: Princeton University Press, 1961.
- Birkhoff, G. Hydrodynamics: A study in logic, fact, and similitude. Second edition. Princeton: Princeton University Press, 1960.
- Birnbaum, A. Some latent trait models and their use in inferring an examinee's ability. In F. M. Lord and M. R. Novick (Eds.), Statistical theories of mental test scores. Reading, Mass.: Addison-Wesley, 1968. Chapters 17-20.
- Bruner, J. S. Some theorems on instruction stated with reference to mathematics. In E. R. Hilgard (Ed.), Theories of learning and instruction. 63rd NSSE Yearbook, Part 1, 1964. Pp. 306-335.

- Bush, R. R., & Mosteller, F. Stochastic models for learning. New York: Wiley, 1955.
- 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.
- Fishman, E. J., Keller, L., & Atkinson, R. C. Massed vs. distributed practice in computerized spelling drills. Journal of Educational Psychology, 1968, 59, 290-296.
- Fries, C. C. Linguistics and reading. New York: Holt, Rinehart, and Winston, 1962.
- Greeno, J. G. Paired-associate learning with massed and distributed repetition of items. Journal of Experimental Psychology, 1964, 67, 286-295.
- Groen, G. J., & Atkinson, R. C. Models for optimizing the learning process. Psychological Bulletin, 1966, 66, 309-320.
- Hartley, R. N. An investigation of list types and cues to facilitate initial reading vocabulary acquisition. Psychonomic Science, 1968, 12, 251-252.
- Hilgard, E. R. (Ed.) Theories of learning and instruction. 63rd NSSE Yearbook, Part 1, 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.
- Laubsch, J. H. An adaptive teaching system for optimal item allocation. Technical Report 151, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1969.

Lorton, P., Jr. Computer-based instruction in spelling: An investigation of optimal strategies for presenting instructional material. U. S. Department of Health, Education, and Welfare. Final Report, Project No. 8-1-026, 1969.

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.: MIT Press, 1968.

Rasch, G. An individualistic approach to item analysis. In P. F. Lazarsfeld and N. W. Henry (Eds.), Readings in mathematical social science. Chicago: Science Research Associates, 1966. Pp. 89-107.

Rothkopf, E. Z. Stimulus similarity and sequence of stimulus presentation in paired-associate learning. Journal of Experimental Psychology, 1958, 56, 114-122.

Smallwood, R. D. The analysis of economic teaching strategies for a simple learning model. Journal of Mathematical Psychology, 1970, in press.

Suppes, P. Some theoretical models for mathematics learning. Journal of Research and Development in Education, 1967, 1, 5-22.

FOOTNOTES

¹An early version of this paper was presented by the first author as an invited address at the Western Psychological Association Meetings, 1969. The second part of the paper was presented at a seminar 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. Support for this research was sponsored by the National Science Foundation, Grant No. NSF-GJ-443X.

²This 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 correct response 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 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 phenomena that the psychologist deals with may be such that it fails to uncover that particular task 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.

