# APPLICATION OF LEARNING MODELS AND OPTIMIZATION THEORY TO PROBLEMS OF INSTRUCTION

by

Verne G. Chant and Richard C. Atkinson

TECHNICAL REPORT NO. 264

October 10, 1975

PSYCHOLOGY AND EDUCATION SERIES

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

TABLE OF CONTENTS

# Application of Learning Models and Optimization Theory to Problems of Instruction*

## I. PHILOSOPHICAL APPROACH

### A. Problem Definition

Stated in its simplest form, the question addressed here is how to allocate instructional resources to achieve a desired objective. Broadly interpreted, this question could include the total educational resources of society and all possible learning situations. In practical terms, however, the setting is restricted to the structural educational system, because this is the only context in which decisions on the allocation of instructional resources may be implemented.

When the question of allocating resources is examined in this setting, attention is usually focussed on a well-defined sub-component of the problem. Once the characteristics of one of these sub-components are understood, their implications may be extended to a larger context. In general, however, the characteristics of many sub-components must be synthesized before solutions can be derived for the problem of resource allocation.

In the school setting, the principal resources to be allocated are the human resources of teachers and students. When the teaching function is augmented by non-human resources, such as computer-aided instruction, then the total instructional resources must be considered. The time spent by the students also must be included because there is frequently a trade-off between instructional resources to be allocated and speed of learning.

There are two basic questions in any resource allocation problem: (1) what are the alternatives and their implications, and (2) which alternative is preferred? The first question concerns the "system" and includes such questions as *what is feasible, what happens if* and *what is the cost?* The second question has to do with the goals, objectives and preferences of the

decision-maker or the collection of people he represents. These are very difficult questions to answer; but they must be answered, at least implicitly, every time an allocation decision is made. This chapter reviews the development and application of mathematical models that help the decision-maker directly with the first question and indirectly with the question of identifying objectives and preferences.

B. Empirical Approach versus Modelling Approach

The core of any decision problem is the determination of the implications or outcomes of each alternative - that is, the determination of the answers to *what happens if?* The questions of feasibility and cost are ancillary to this central problem and are relatively uncomplicated. For example, consider the problem of determining optimal class size. For a particular situation, the question of feasibility might involve simply the availability of physical facilities and instructional resources. Analysis of the question of cost also would be reasonably straightforward. It would, however, be very difficult to determine and quantify the expected results with sufficient accuracy to permit assessment of the cost-effective trade-off. It is the quantitative analysis of the core of the decision problem that can be approached with empirical or modelling techniques.

In the empirical approach, the input variables (class size, for example) and the output variables (amount learned, say) are defined for the particular problem at hand and then empirical data relating to these variables are collected and analyzed. From the analysis it is hoped that a causal relationship can be determined and quantified. This relationship then serves to predict the output from the system for the range of alternatives under consideration. Once the expected output has been quantified and once the costs of the alternatives have been determined, the decision problem is reduced to an evaluation of preferences.

The empirical approach has a natural appeal for several reasons. First, perhaps, is its simplicity. If a particular system has only a few variables that are amenable to quantification, then, given sufficient data, the relationships between them can be determined. The second reason for its appeal is that no *a priori* knowledge of the relationships among variables is necessary; the data simply speak for themselves. A third reason is that data analysis can never really be avoided completely, whatever approach is employed. Thus, if the problems of data collection, verification and analysis must be encountered regardless, it may appear expeditious to rely on data analysis alone.

There are, however, many problems with the application of the empirical approach, especially to situations that are as complicated as those that comprise the educational system. It is extremely difficult to define real variables precisely. Often surrogate variables must be used because the real variables cannot be suitably quantified. For example, teaching ability can be represented by such quantifiable variables as years-of-experience and level-of-education. Even if variables can be defined, the complexities of measurement introduce new problems. These problems involve statistical sampling, measurement error and the choice of survey and interview techniques.

In addition to definitional and measurement problems, difficulties arise in controlling multiple variables and long time constants or reaction times. Within a system of many variables, the relationships between only a few of them may be impossible to extract empirically because of the influence of other uncontrolled or unquantified variables. Moreover, the fact that educational systems have long time constants introduces complications when more than "snapshot" data analysis is required. Time series or

"longitudinal" data analysis is particularly important when the objective is to study the effects resulting from a change in the system, whether it be an experimental change or a permanent change. Because of the long time constants in education, the effects of change are manifested very slowly and the detection of the change through data analysis requires the maintenance of high quality data over a relatively long time period.

The second method of analyzing the system is the modelling approach. This approach is characterized by some assumptions about the structure of the system - that is, it assumes a particular form for relationships among some of the variables. It encompasses a spectrum of techniques ranging from structured data analysis to abstract theory.

In its most abstract form, the modelling approach offers the power of mathematical analysis with the capability of examining a wide range of alternatives or parameter values. The models that result from fitting mathematical equations to empirical data also may be amenable to mathematical analysis; but often, because of their complexity, they require the power of computers to analyze the effects of various alternatives and parameter values. It is, of course, possible to combine the abstract model form with extensive data analysis. Indeed, the optimal balance of model abstraction and data analysis is the goal of any model builder. This balance depends upon many factors, including the purpose of the model, the availability of appropriate data and the characteristics of the decision-maker as well as the analyst. A good model is characterized by providing sufficient detail for the decision-maker while retaining no more complexity than is required to portray adequately relationships within the real environment.

C. Mathematical Models and Optimization Theory

A particularly useful form of the modelling approach is one in which the problem is formulated within the framework of control and optimization theory. At the heart of this framework is the mathematical model that is a dynamic description of the fundamental variables of the system. For any alternative under consideration, the model determines all the implications or outcomes over time resulting from the implementation of that particular alternative or policy.

Once the implications of each alternative are known and the costs have been evaluated, preferences can be assigned to the various alternatives. In the framework of control and optimization theory, these alternatives for resource allocation are associated with settings of the control variables. The preferences over all possible alternatives are specified by an objective function that measures the trade-off between benefits and costs, which are defined in the model by the values of the control variables and the state variables. The control and state variables define, generally speaking, the inputs and outcomes of a system, respectively. The problem of optimal resource allocation is thus the problem of choosing feasible control variable settings that maximize (or minimize) the objective function.

The central dynamic behavior that must be modelled when considering problems of resource allocation in the educational setting is the interaction between the instructor - whether it be teacher, computer-assisted instruction or programmed instruction - and the individual learner. The effects of the environment (for example, the classroom) also are important. Models of these interactions are essential in order to predict the outcomes of alternative instructional policies. Once the cost components of the various alternatives have been evaluated, the optimization problem may take one of three forms.

If the quantity of resources is fixed, then benefits can be maximized subject to this resource constraint. If there is a minimum level of performance to be achieved, then the appropriate objective is to minimize cost subject to this performance level. Finally, if performance and cost are both flexible and if the trade-off of benefit and cost can be quantified in an objective function, then both the optimal quantity of resources and the level of performance can be determined.

## II. PREVIOUS RESEARCH

### A. Overview

The applications of learning models and optimization theory to problems of instruction fall into two categories: (1) individual learner oriented, and (2) group of learners (classroom) oriented. In category (1) applications, instruction is given to one learner completely independently of other learners. These applications are typical of computer-assisted instruction and programmed instruction and also include the one-teacher/one-student situation. Within this category, many situations can be adequately described by an appropriate existing model from mathematical learning theory. In such cases, as outlined below, the results of applying mathematical models have been encouraging. In other more complex situations, existing models must be modified or new models must be developed to describe the instructor/learner interaction.

In category (2) applications, instruction is given simultaneously to two or more learners. This characteristic is typical of classroom-oriented instruction and also includes other forms of instruction, such as films and mass media, where two or more learners may be receiving instruction but there is no feedback from learner to instructor. In contrast to category (1) situations, where mathematical learning theory provides suitable models of instructor/

learner interaction, there is no comparable theory for the group of learners environment. Category (2) applications must therefore include model development as well as mathematical analysis.

Most applications, whether in category (1) or in category (2), follow a four-step procedure.

Step one is to isolate a particular learning situation. In this step, the learning situation is classified as category (1) or (2), the method of instruction is defined and the material to be learned is specified.

Step two is to acquire a suitable model to describe how instruction affects learning. This step may be as simple as the selection of an appropriate model from mathematical learning theory, as mentioned above, or as difficult as the development of a new model for the particular situation.

Step three is to define an appropriate criterion for comparing the various instruction possibilities, taking account of benefits and costs as determined by the model.

Step four is to perform the optimization and analyze the characteristics of the optimal solution. These characteristics may include the sensitivity of the optimal solution to key variables of the model and the comparison of its results relative to those of other solutions. In some situations the optimization problem may be very difficult or impossible to solve. In this case, various sub-optimal solutions may be identified whose results represent improvements over those of previous solutions.

B. Individual Learner Setting

1. *Quantitative Approach for Automated Teaching Devices*

An important application of mathematical modelling and optimization theory was the development of a decision structure for teaching machines by Smallwood (1962). Smallwood's goal was to produce a framework for the

design of teaching machines that would emulate the two most important qualities of a good human tutor: (1) the ability to adjust instruction to the advantage of the learner, and (2) the ability to adapt instruction based on his own experience. The decision system within this framework must therefore make use of the learner's response history, not only to the benefit of the current learner, but also for future learners.

The learning situation considered by Smallwood has three basic elements: (1) an ordered set of concepts that are to be taught, (2) a set of test questions for each concept to measure the learner's understanding, and (3) an array of blocks of material that may be presented to teach the concepts. Two additional elements are required to complete the framework for the design of a teaching machine: (4) a model with which to estimate the probability that a learner with a particular response history will respond with a particular answer to each question, and (5) a criterion for choosing which block to present to a learner at any given time.

Having defined his model requirements in probabilistic terms, Smallwood considered three modelling approaches: correlation, Bayesian and intuition. He discarded the correlation model approach as not useful in this context. Then he developed Bayesian models, based on the techniques of maximum likelihood and Bayesian estimation (these models are too complex to review here). His intuition approach led to a relatively simple quantitative model based on four desired properties: representation of question difficulty and learner ability, together with model simplicity and experimental performance.

The model is

$$P = \begin{cases} \frac{bc}{a} & b \leq a \\ 1 - \frac{(1-b)\ (1-c)}{(1-a)} & b > a \end{cases}$$

where $P$ is the probability of a correct response, $b$ measures the ability of the learner, $c$ measures the difficulty of the question and $a$ is an average of the fraction of correct responses. All parameters are between zero and one.

As an objective function for determining optimal block presentation strategies, Smallwood suggested two possibilities with variations. One was an amount-learned criterion, which measured the difference before and after instruction, and the other was a learning-rate criterion, which essentially normalized the first criterion over time. In the optimization process, these criteria are used to choose among alternative blocks for presentation in a local, rather than global sense.

A simple teaching machine was constructed based on the concepts of this decision structure. The experimental evidence verified that the machine distinguished between learners and presented them with different combinations of blocks of material. It also verified that different decisions were taken at different times under similar circumstances, indicating that the machine was adaptive.

2. *Order of Presentation of Items from a List*

The task of learning a list of paired-associate items has practical applications in many areas of education, notably in reading and foreign language instruction (Atkinson, 1972). It is also a learning task for which models of mathematical learning theory have been very successful at describing empirical data. Is is therefore not surprising that the earliest and most encouraging results of the application of optimization techniques have come in this area. Although the learning models employed in these studies are extremely

simple, the results are valuable for three reasons: (1) the applications are practical, (2) these results lead to further critical assessment of the basic learning models, and (3) the general analytical procedure is transferable to more complex situations.

The application of mathematical models and optimization theory to the problem of presenting items from a list can be illustrated by three examples from the literature. The first is a short paper by Crothers (1965) that derives an optimal order of item presentation when two modes of presentation are available. The second is an in-depth study by Karush and Dear (1966) of a simple learning model that leads to an important decomposition result. The third example is a paper by Atkinson and Paulson (1972) that derives optimal presentational strategies from three different learning models and presents some experimental results. These three papers are described briefly.

In the Crothers paper there are two modes of presentation of the items from the list; the total number of presentations using each mode is fixed, but the order of presentation is to be chosen. Since the order of presentation does not affect the cost of the instruction, the objective is simply to maximize the expected proportion of correct items on a test after all presentations have been made.

Two models of the learning process are studied in this paper. The *random trial increment* model (which is described in detail later in this section) predicts that the expected proportion of correct items is independent of the order of presentation of items; therefore, any order is an optimal solution. The second learning model, the *long-short learning and retention* model, predicts different results from different presentation orders, and so a meaningful application of optimization exists. This model depicts the learner

as being in one of three states: a learned state, a partial learning state and an unlearned state. The learner responds with a correct respose with probability *1*, *p* or *g*, respectively, depending upon his state of learning, and his transition from state to state is defined by the probabilistic transition matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ a & 1-a & 0 \\ b & c & 1-b-c \end{bmatrix} .$$

This model simplifies into the *two-element* model by setting $b$ equal to zero and further into the *all-or-none* model by dropping the partial learning state. This model is assumed to describe the learning process for each mode of presentation, so that the response probabilities for each state are identical for all modes but the parameters $a$, $b$ and $c$ are different for each mode. For a discussion of these models, see Atkinson, Bower and Crothers (1965).

The result of the optimization step in this application is contained in two theorems. The first theorem states that the ranking of presentation schedules based on the expected proportion of correct responses (which is the defined objective) is identical to the ranking based on the probability of occupying the learned state. The second theorem states that the ranking of two presentation schedules is preserved if the schedules are either prefixed or suffixed by identical strings of presentations. These theorems are sufficient to conclude that moving one presentation mode to the right of another in a schedule always has the same (qualitative) effect on the terminal proportion correct and, hence, that optimal presentation schedules have all presentations of one mode together.

In the learning situation described by Karush and Dear, there are

$n$ items of equal difficulty to be learned, and the problem is to determine which item out of the $n$ to present for study at any given time. The strategy for choosing items for presentation is to take into account the learner's response history up to the current time. The all-or-none model is used to describe the learning process, and it is assumed that the single model parameter has the same value for each item.

In order to formulate an objective function, it is assumed that all presentational strategies have the same cost so that the objective can be defined in terms of the state of learning at the termination of the strategy. Assuming that all items are weighted equally, an expected loss function is defined in terms of the probabilities $P_k$ that at the terminal node exactly $k$ items are still unlearned. The expected loss for a particular terminal node is given by

$$\sum_{k=o}^{n} P_k b_k$$

where $b_k$ is the value (weight) of the loss if $k$ items are still unlearned. The overall expected loss, which is to be minimized, is therefore

$$\sum_{h} q(h) \quad \sum_{k} P_k(h) b_k$$

where $q(h)$ is the probability of occupying terminal node $h$ and the first summation is over all possible terminal nodes. For the particular values $b_k = 1$, the objective function above is equivalent to the maximization of the probability that all items are learned; and for $b_k = k$ it is equivalent to the maximization of the expected sum of the probabilities of being in the learned state for each item. All of the results that are derived in the paper are

not dependent on the values for the $b_k$, and so they are quite general.

The optimization is accomplished using the recursive formulation of dynamic programming. The principal result is that, for arbitrary initial probabilities of being in the learned state for each item, an optimal strategy is to present the item for which the current probability of learning is the least. The most practical application of the results is for the case where these initial probabilities are zero, in which case the optimal strategy can be implemented simply by maintaining counts of correct and incorrect responses on each item. Also in this case, the optimal strategy is independent of both model parameters: the probability of transition and the probability of guessing.

Atkinson and Paulson report empirical results employing the all-or-none based optimal strategy derived by Karush and Dear and compared it with strategies based on other learning models. In one experiment, the all-or-none-based strategy is compared with the optimal strategy derived from the linear model. In the derivation of this latter optimal strategy, it is assumed that the model parameters are identical for all items. For the objective of maximizing the expected number of correct responses at the termination of the experiment, it is shown that all items should be presented the same number of times. Consequently, a random-order strategy is employed in which all items are presented once, then randomly reordered for the next presentation and so on. The experimental results show that during the learning experience the all-or-none-based strategy produces a lower proportion of correct responses than the linear-based (random) strategy, but that on two separate post-experiment tests, the all-or-none-based strategy yields a higher proportion of correct responses. From these results it can be concluded that in this learning situation and for the stated objective the all-or-none model described the learning process more accurately than the linear model.

In another experiment, the all-or-none-based strategy and the linear-based strategy are compared with a strategy based on the random trial increment (RTI) model. The RTI model is a compromise between the all-or-none and the linear models. Defined in terms of the probability $p$ of an error response, at trial $n$ this probability changes from $p(n)$ to $p(n + 1)$ according to

$$p(n + 1) = \begin{cases} p(n) \text{ with probability } 1 - c \\ ap(n) \text{ with probability } c \end{cases}$$

where $a$ is a parameter between zero and one and $c$ is a parameter that measures the probability that an "increment" of learning takes place on any trial. This model reduces to the all-or-none model if $a = 0$ or to the linear model if $c = 1$.

This application of the RTI model differs in two ways from the earlier studies outlined above. First, because of the complexity of the optimization problem, only an approximation to the optimal strategy is used. The items to be presented at any particular session are chosen to maximize the gain on that session only, rather than to analyze all possible future occurrences in the learning encounter. Second, the parameters of the model are not assumed to be the same for all times. These parameters are estimated in a sequential manner, as described in the Atkinson and Paulson paper; as the experiment progresses and more data become available regarding the relative difficulty of learning each item, refined estimates of the parameter values are calculated.

The results of the experiment show that the RTI-based strategy produces a higher proportion of correct responses on post-tests than either the all-or-none-based or linear-based strategies. The favorable results are due partly to the more complex model and partly to the parameter

differences for each item. This conclusion is supported by the fact that the relative performance of the RTI-based strategy improves with successive groups of learners as better estimates of the item-related parameters are calculated.

3. *Interrelated Learning Material*

In many learning environments, the amount of material that has been mastered in one area of study affects the learning rate in another distinct but related area - for example, the curriculum subjects of mathematics and engineering. In situations such as this, the material in two related areas may be equally important, and the problem is to allocate instructional resources in such a way that the maximum amount is learned in both areas. In other situations, the material in one area may be a prerequisite for learning in another rather than a goal in itself. Here, even though the objective may be to maximize the amount of material learned in just one area, it may be advantageous in the long run to allocate some instructional resources to the related area. This problem of allocating instructional effort to interrelated areas of learning has been studied by Chant and Atkinson (1973). In this application, a mathematical model of the learning process did not exist, and so one had to be developed before optimization theory could be applied.

The learning experience from which the model was developed was a computer-assisted instructional program for teaching reading (Atkinson, 1974). This program involved two basic interrelated areas (called strands) of reading, one devoted to instruction in sight-word identification and the other to instruction in phonics. It has been observed that the instantaneous

learning rate on one strand depended on the student's position on the other strand.

In the development of the learning model, it was assumed that the interdependence of the two strands was such that the instantaneous learning rate on either strand is a function of the difference in achievement levels on both strands. Typical learning rate characteristics are shown in Figure 1. If the achievement levels on the two strands at time $t$ are represented by $x_1(t)$ and $x_2(t)$, then the instantaneous learning rates are the derivatives of $x_1$ and $x_2$ with respect to time; these rates are denoted as $\dot{x}_1$ and $\dot{x}_2$. By defining $u(t)$ as the relative amount of instructional effort allocated to strand one, the model of learning can be expressed in differential equation form as

INSERT FIG 1

$$\dot{x}_1(t) = u(t) f_1(x_1(t) - x_2(t)),$$

$$\dot{x}_2(t) = [1 - u(t)]\, f_2(x_1(t) - x_2(t)),$$

where $f_1$ and $f_2$ are the learning rate characteristic functions depicted in Figure 1. In this formulation of the problem, the total time, $T$, of the learning encounter is fixed and the objective is to maximize a weighted sum of the achievement levels on the two strands at the termination of the encounter. The objective is therefore to maximize

$$c_1 x_1(T) + c_2 x_2(T),$$

where $c_1$ and $c_2$ are given non-negative weights. This maximization is with respect to $u$ subject to the constraint $0 \leq u(t) \leq 1$ for all $t$ such that $0 \leq t \leq T$.

The optimization is carried out, not for the nonlinear learning rate characteristic functions of Figure 1, but for linearized approximations

to them. From the form of the optimal solutions, it is clear that the analysis applies equally well to the nonlinear functions. The optimization is performed by means of the Pontryagin Maximum Principle. It is shown that the optimal solution is characterized by a "turnpike" path in the $x_1x_2$ plane. On the turnpike path the difference $x_1 - x_2$ between the achievements levels on the two strands remains constant. Optimal trajectories are such that initially all of the instructional effort is allocated to one of the strands until the turnpike path is reached. Then the instructional effort is apportioned so as to maintain a constant difference between strands - that is, so as to remain on the turnpike path. Near the end of the learning encounter, the instructional effort is again allocated to just one strand, depending on the relative values of the weights $c_1$ and $c_2$ of the objective function. Figure 2 shows the turnpike path and typical optimal trajectories starting from two different initial points and terminating according to two different values of objective function weights.

INSERT FIG 2

It is also shown that of all the stable paths, the turnpike path is the one on which the average learning rate is maximized. A stable path is the steady state path that is approached if the relative allocation of instructional effort between strands is held constant. It can be shown that stable paths are such that the difference between achievement levels on the two strands is constant.

C. Group of Learners Setting

1. *A Descriptive Model Structure*

Carroll (1965) developed a structure for describing learning in the school or classroom setting. This model involves five variables; four are defined in a quantitative sense, but one is difficult to quantify. The relationships among these variables are not precisely defined, but the potential interactions are identified and described.

The five variables are aptitude, perseverance, ability to comprehend instruction, quality of instruction and opportunity to learn. The aptitude variable is defined as a reference learning rate for a learner for a given task. Aptitude is to be measured by the reciprocal of the time required to master the given task to a given criterion under optimal learning conditions. The perseverance variable is defined by the length of time that the learner is willing to spend learning the task involved. Carroll suggests that this variable will change significantly over time and that it can be affected by external factors. The variable ability-to-comprehend-instruction is assumed to be primarily represented by verbal intelligence, and so measures of verbal intelligence are considered adequate for quantification purposes. It is suggested that this variable will demonstrate less rapid changes over time than, for example, perseverance and that it is determined to a large extent by the individual's early life environment. Carroll's fourth variable, quality of instruction, is defined imprecisely as the degree to which content and method of instruction are structured so that material is easily learned. There is an important joint relationship between quality

of instruction and ability to comprehend instruction on the learning rate. This relationship is such that low quality instruction more severely hinders the learner with limited ability to comprehend instruction that the learner with greater ability. The final variable, opportunity to learn, is defined as the time actually allowed for learning in the particular situation. It is recognized that in the classroom not all learners have a continuous opportunity to learn since the class must learn together.

Without more explicit elaboration of the relationships among these variables, and in some cases more precise definitions, this model cannot be used in a quantitative sense. It has been very useful, nevertheless, to help identify the salient features of the learning process in the classroom.

2. *Normative Models*

Restle (1964) made an early contribution to the application of learning models and optimization theory to the classroom or group of learners setting. He has studied two situations, each of which involves a group of identical learners. In the first situation, the problem is to determine the optimal class size for a large number of identical learners. The objective function is expressed in cost terms, including both instructor and learner costs, and the amount to be learned is fixed. In the second situation, the problem is to determine the optimal pace of instruction for a curriculum consisting of a sequence of identical items in which further learning progress for any learner is terminated if an item is not mastered. The pace of instruction is determined by the amount of time allocated to each item, assuming equal time for each item and a fixed

total amount of time. The objective is to maximize the expected number of items learned by the group or, equivalently, by any learner of the group.

The continuous time all-or-none model is used to describe the learning process in both situations. This version of the all-or-none model is essentially the same as the discrete (learning trial) version introduced earlier and is defined by the cumulative distribution function

$$F(t) = 1 - e^{-\lambda t}$$

which gives the probability that learning on an item takes place before time $t$, where $\lambda$ is the reciprocal of the mean time until learning occurs.

For the optimal class size situation, Restle chooses to minimize the expected total (weighted) time cost of both instructors and learners, subject to the constraint that instruction be given until all learners have mastered the item. Based on the model, the expected time $M(n)$ for a group of $n$ learners to learn an item is given by

$$M(n) = \frac{1}{\lambda} \sum_{k=1}^{n} \frac{1}{k} .$$

Letting $r$ represent the ratio of the value of instructor time to the value of learner time leads to the expression

$$NM(n) + rNM(n)/n$$

for expected total time cost in learner time units where $N$ is the total number of learners and $n$ the size of each sub-group (assuming that $N$ is large enough that the integrality error is negligible). Using a continuous approximation for $M(n)$, this optimization is easily performed to yield the

relationship shown in Figure 3 between optimal class size and $r$, the relative value of instructor and learner time.

INSERT FIG 3

For the situation involving optimal pace of instruction, the total amount of time $(T)$ is allocated equally to each item in order to maximize the expected number of items mastered by a learner. If $t$ units are allocated to each unit, then, based on the model, the mean number of items learned is

$$e^{\lambda t} - 1 - e^{\lambda t}(1 - e^{-\lambda t})(T + t)/t$$

Rather than calculate the maximum of this expression with respect to $t$, Restle shows the function graphically for various values of the basic parameter $T\lambda$. With this learning model, $T\lambda$ represents the expected number of items learned for an individual learner who is allowed to proceed to the next item as soon as he has mastered the current item. On the basis of the graphs, Restle concludes that for a short course where $T\lambda = 3$, the optimal pace for a group is instruction on 2 items. For a medium-length course of $T\lambda = 12$, the group should receive instruction on 4 items; and for a long course with $T\lambda = 144$, the group takes 30 items. Thus, for long sequences of items in which a learner is blocked if he misses only one item, the group pace must be very slow compared to the tutored pace.

In a paper by Chant and Luenberger (1973), a mathematical theory of instruction has been developed that describes certain aspects of the classroom environment. This model is developed in two stages; the first models the instructor/learner interaction for an individual

learner situation, and the second extends this model to a group of learners situation. In the first stage, the principal problem under investigation is the optimal matching of instruction to the characteristics of the learner. In the second stage, the analysis is concerned with the problem of instruction pacing, which is an important question in the classroom situation.

Motivated by a differential equation formulation of the learning curve by Thurstone (1930), Chant and Luenberger assume that the relationship between learning rate, instructional input and state of the learner can be represented by

$$\dot{p}(t) = u(t)g(p(t))$$

where $p(t)$ is the achievement level of the learner at time $t$ relative to total learning. In this equation $\dot{p}(t)$ represents learning rate, $u(t)$ is an instructional input variable and $g$, the *characteristic learning function*, describes how learning rate depends on the achievement level for a particular learner in a particular situation. Restrictions are placed on the function $g$, so that for a constant instructional input $u(t)$ the learning curve has the familiar S-shape.

The instructional input variable $u(t)$ is thought of as a measure of the *intensity of instruction* in the sense that the larger the value of $u(t)$, the greater the learning rate and the cost of instruction. The relationship between instruction cost and learning rate (for a given achievement level) forms the basis of the precise definition of $u(t)$ such that the total cost of instruction for $t = 0$ to $t = T$ is

$$\int_0^T (u(t))dt$$

where $\ell(u(t))$ defines the rate of expenditure of instructional resources for instruction of intensity $u(t)$, $0 \leq t \leq T$.

In formulating an objective function, both the learner's achievement level and the cost of the learning encounter are considered. The learner's achievement level at the end of the encounter is represented by $p(T)$ and the cost of the learner's time by $bT$. The objective function is defined as the net benefits; that is

$$p(T) - bT - \int_0^T \ell(u(t))dt.$$

The relative importance of achievement level and instruction cost is assumed to be included in the loss function $\ell$.

The optimization problem is to choose the instructional input $u(t)$ for $0 \leq t \leq T$ and the duration $T$ of the learning encounter so as to maximize the above objective function. It is shown in the paper that the optimal instructional input function $u$ is constant throughout the learning encounter and is determined by the solution of

$$u\ell'(u) - \ell(u) - b = 0.$$

The optimal value of $T$ is given by the larger of the two values that safisfy

$$g(p(T)) = \ell'(u).$$

The result that the optimal instructional input is constant throughout the learning encounter is quite general in that it does not depend on

the particular characteristic learning function or the particular loss function.

In the second stage of their development of a mathematical theory of instruction, Chant and Luenberger first define a learner aptitude parameter that is used to characterize the diverse nature of a nonhomogeneous group of learners. Aptitude is defined in a relative sense by comparing the learning times of two learners under identical situations. One learner is said to have an aptitude twice as great as another if he learns the same amount in half the time. This definition is similar to Carroll's mentioned above. Using this concept of aptitude, the characteristic learning function $g$ is redefined such that the aptitude component is separated from the other components. The basic instructor/learner model now becomes

$$\dot{p}(t) = u(t)ag(p(t)).$$

The above optimization is unchanged with this modification, so that the optimal instructional input is still constant over time.

The development of the group learning model for the purpose of determining the optimal pace begins with an analysis of the relationship between pace and aptitude for an individual learner. To model the effect of instruction pacing, a body of sequential learning material is divided into a sequence of blocks. The basic instructor/learner model outlined above is used to describe the learning process on each block. The sequential nature of the material is captured by specifying how the learner's performance on one block depends on his achievement on preceding blocks. This interblock dependence is defined by the *block inter-*

*action function* $h$, which relates the initial state on a block to the final achievement level on the preceding block. For analytical purposes, an infinite sequence of similar blocks is considered. Blocks are similar if the learning performance for them can be described with identical characteristic learning functions and block interaction functions. The infinite sequence is considered in order to eliminate transient effects and to concentrate on steady state relationships. An infinite sequence of similar blocks is illustrated in Figure 4.

INSERT FIG 4

The steady state learning behavior of a learner on an infinite sequence of similar blocks is characterized by allocating an equal amount of instructional time to each block and determining the achievement level that the learner approaches on each block as the number of blocks increases towards infinity. The *pace* of instruction is defined as the amount of time $\tau$ that is spent on each block. In the limit, the initial state on each block is the same, the final achievement level on each block is the same and the pace is such that the learner progresses from this initial state to this final level. This steady state condition is illustrated in Figure 5.

INSERT FIG 5

For an individual learner with a particular S-shaped learning curve and block interaction function, the correspondence between pacing $\tau$ and the steady state final achievement level is defined as the *steady state response function* $p_s$. With suitable assumptions, it can be shown that $p_s(\tau)$ is zero for $\tau < \tau_c$, where $\tau_c$ is defined as the *critical pace*, that $p_s$ is concave and increasing for $\tau > \tau_c$ and has infinite slope at $\tau = \tau_c$.

For determining the optimal pace of instruction, the objective function of steady state achievement level on a block per unit of time on the block is defined. This ratio, called *gain* and denoted $\gamma$, is given by

$$\gamma(\tau) = p_s(\tau)/\tau.$$

The maximization of gain implies that

$$p_s(\tau) = \tau p_s'(\tau).$$

This relationship is illustrated in Figure 6.

INSERT FIG 6

The *steady state response reference function* $p_r$ is defined as the function $p_s$ but for a learner with unity aptitude. In view of the definition of aptitude as the reciprocal of learning time, the response of a learner with aptitude $a$ for pacing $\tau$ is simply $p_r(a\tau)$.

A nonhomogeneous group of learners is characterized by the aptitudes of the learners in the group with the assumption that all the learners have identical characteristic learning functions and block interaction functions. The objective function for the group, called *group gain* and denoted $\Gamma$, is defined by

$$\Gamma(\tau) = (1/\tau) \sum_{i=1}^{N} p_r(a^i \tau)$$

where the $N$ learners of the group have aptitudes $a^i$, $i = 1$ to $N$. The optimal group pace is defined by the maximization of this group gain. It is shown that for widely diverse groups, the optimal pace is such that the lower aptitude part of the group has a zero steady state response; that is, these learners are dropped from the group because of the fast pace. In addition, for homogeneous groups, the optimal group pace is the same as the optimal individual learner pace for that aptitude.

## III. AREAS OF FURTHER RESEARCH

This concluding section is intended to highlight a few areas in the field of application of learning models to problems of instruction that require further work. In addition, suggestions are given as to the research directions that may be most effective for making these applications more practical.

### A. Problems of Measurement

Problems of measurement exist when we cannot quantify exactly what we want quantified. In order to verify a quantitative model empirically or to apply it in real world situations, the variables of the model must be measurable. The measurement process can be complicated at either of two levels: the variables of the model may not be satisfactorily quantifiable or, if quantifiable, there may be estimation problems - that is, there may be no satisfactory method of determining a unique value for the defined variable.

To illustrate these two kinds of problems, consider a situation where it is required to have a measurement on the state of a learner with respect to some set of material. At the outset, the first kind of problem is evident since a precise definition of the variable concerned is not available. A satisfactory solution to this problem is perhaps to define a surrogate variable that represents the real variable. In this situation, a proportional measure of the learner's knowledge of the material as indicated by his score on some test may be

an adequate surrogate variable. The second kind of problem has to do with the variability of tests themselves and the learner's performance on them. Different tests that are intended to measure equivalently the set of material involved will yield different results and the results on a particular test are affected by the testing environment, by guessing and by numerous other factors.

In experimental situations, these problems can be alleviated to a certain extent by careful design. In these situations, the set of material that is to be learned is chosen so that it may be described precisely and simply - for example, in paired-associate learning experiments. This simplifies both the knowledge definition problem as well as the estimation problem. In real applications, however, these problems can be severe.

These problems of measurement can best be attacked during the formulation and modelling phases of the analysis of problems of instruction. It is of limited use to have a model that cannot be investigated experimentally. It is of no practical value to have an experimentally verified instructional technique that requires such extensive measurement and data analysis that implementation is not cost effective. These measurement problems must be considered during the overall analysis. In some cases, they may be alleviated at implementation by having an estimation model incorporated as part of the technique to be applied.

B. Individual Learner Setting

Optimizing the performance of individual learners is an area that has tremendous potential for impact, even though it has already

received some attention. The application of mathematical models and optimization theory to learning problems in computer-aided instruction is likely to prove increasingly useful in the future. Complex models of learning must be developed, and they should be designed for implementation in particular situations. These models have to be complex so as to describe adequately the particular learning phenomena in the situation; but such complexity is manageable provided that the models can be adapted for computer implementation. What is needed, then, is a clear understanding of the ultimate application of the model so that its development is guided by the requirements of implementation.

C. Classroom Setting

Developments in the classroom setting are much farther from implementation than those for individual learner setting. For the classroom, general models must be developed that cover broad categories of learning and instruction. Existing models must be extended and new models must be developed to account for group learning phenomena that so far have been ignored or not even identified. To accomplish this, theoretical and empirical research must complement and supplement each other. Similarly, work by researchers in education and psychology must be continually synthesized. One promising avenue to pursue in this respect would be to engage in model-directed data analysis; that is, either by using an existing model or by developing a model appropriate for the situation to be investigated, data gathering experiments and analyses should be designed and carried out to verify or refute these

models. In this approach, the model directs the empirical research by imposing a structure on the system or by proposing relationships or conclusions to be tested. In this way, the complex relationships that comprise an educational system can be more readily isolated, and hence more easily understood.

REFERENCES

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

Atkinson, R.C. Teaching children to read using a computer. American Psychologist, 1974, 29, 169-178.

Atkinson, R.C., Bower, G.H. and Crothers, E.J. An introduction to mathematical learning theory. New York: Wiley, 1965.

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

Carroll, John B. School learning over the long haul. In Krumboltz, J.D., (Ed.), Learning and the educational process, Chicago: Rand McNally, 1965.

Chant, V.G. and Atkinson, R.C. Optimal allocation of instructional effort to interrelated learning strands. Journal of Mathematical Psycology, 1973, 10, 1-25.

Chant, V.G. and Luenberger, D.G. A mathematical theory of instruction: Instructor/learner interaction and instruction pacing. Journal of Mathematical Psychology, 1974, 11, 132-158.

Crothers, E.J. Learning model solution to a problem in constrained optimization. Journal of Mathematical Psychology, 1965, 2, 19-25.

Karush, W. and Dear, R.E. Optimal stimulus presentation strategy for a stimulus sampling model of learning. Journal of Mathematical Psychology, 1966, 3, 19-47.

Restle, F. The relevance of mathematical models for education. In Hilgard, E.R. (Ed.), Theories of learning and instruction: 63rd yearbook of the national society for the study of education, Chicago: University of Chicago Press, 1964.

Smallwood, R.D. A decision structure for teaching machines. Cambridge: MIT Press, 1962.

Thurstone, L.L. The learning function, Journal of General Psychology, 1930, 3, 469-493.

FOOTNOTE

*Preparation of this paper was supported in part by Office of Naval Research Contract N00014-67-A-0012-0054 and a grant from the National Institute of Mental Health (MH-21747).

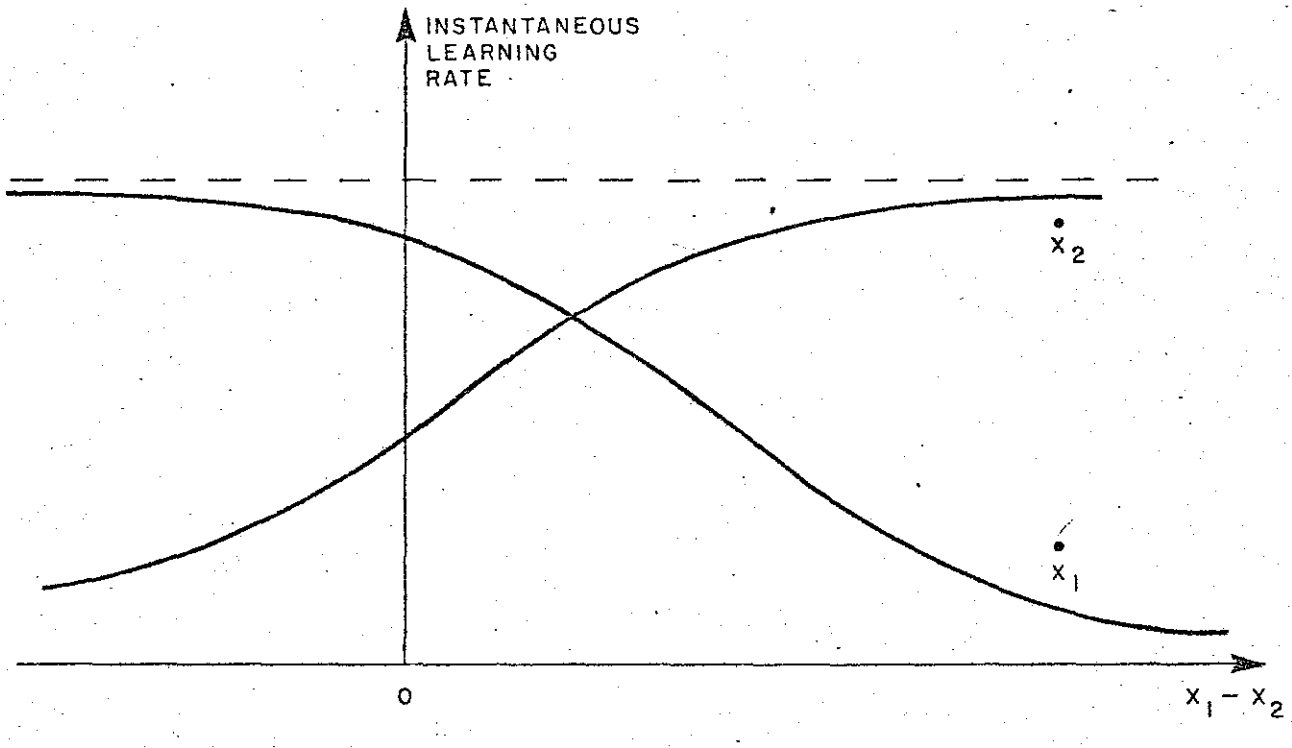


Figure 1. Typical learning rate characteristics.

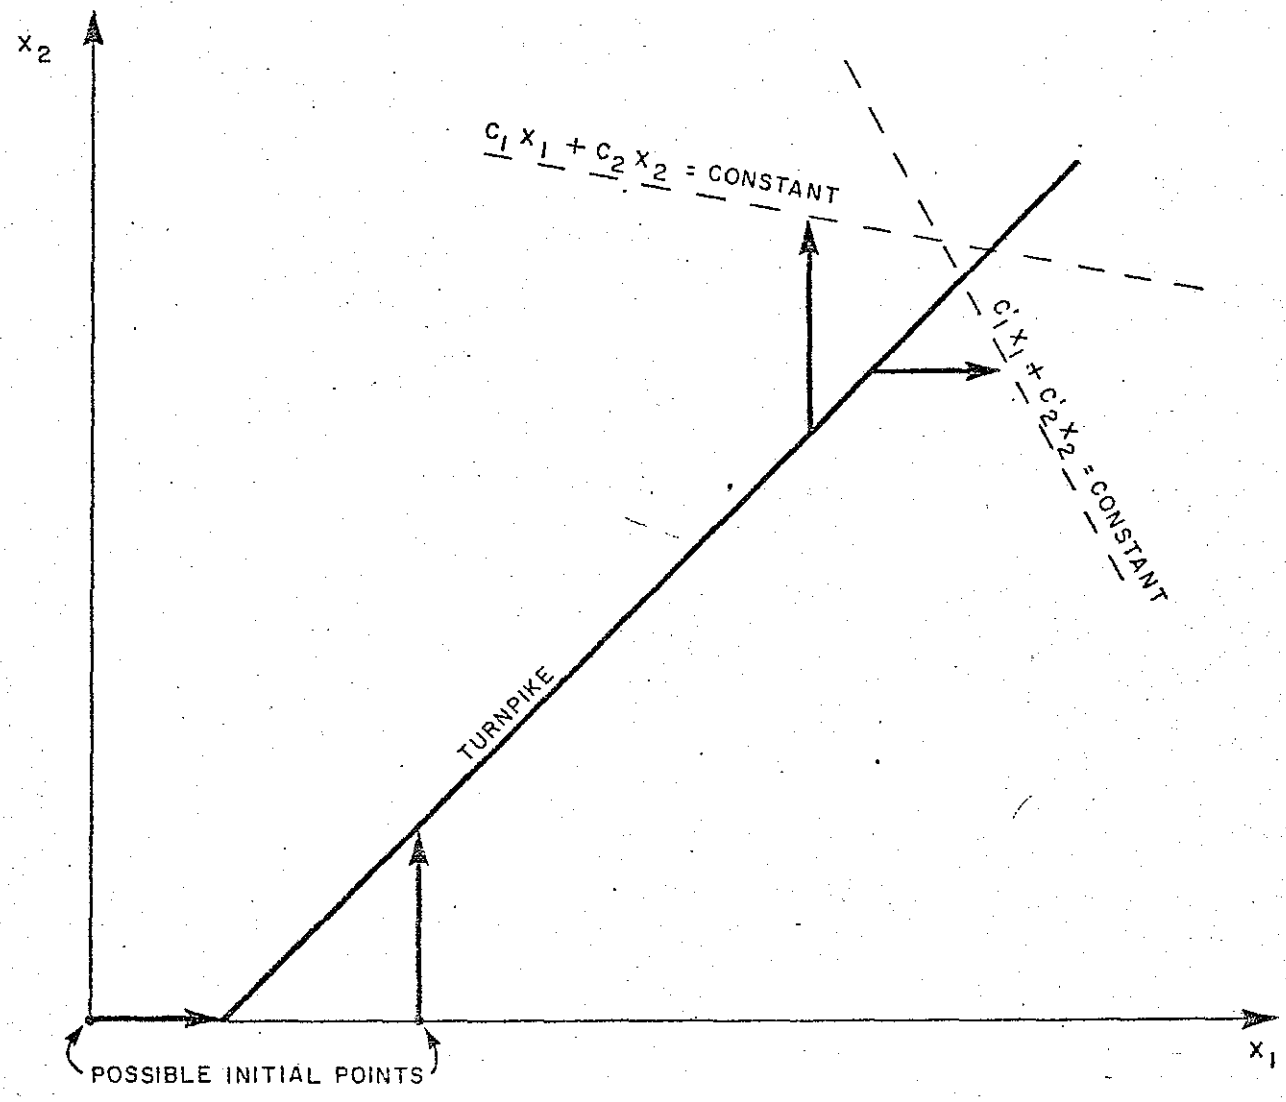


Figure 2. Optimal trajectories using turnpike from two possible initial points and with two possible objectives.

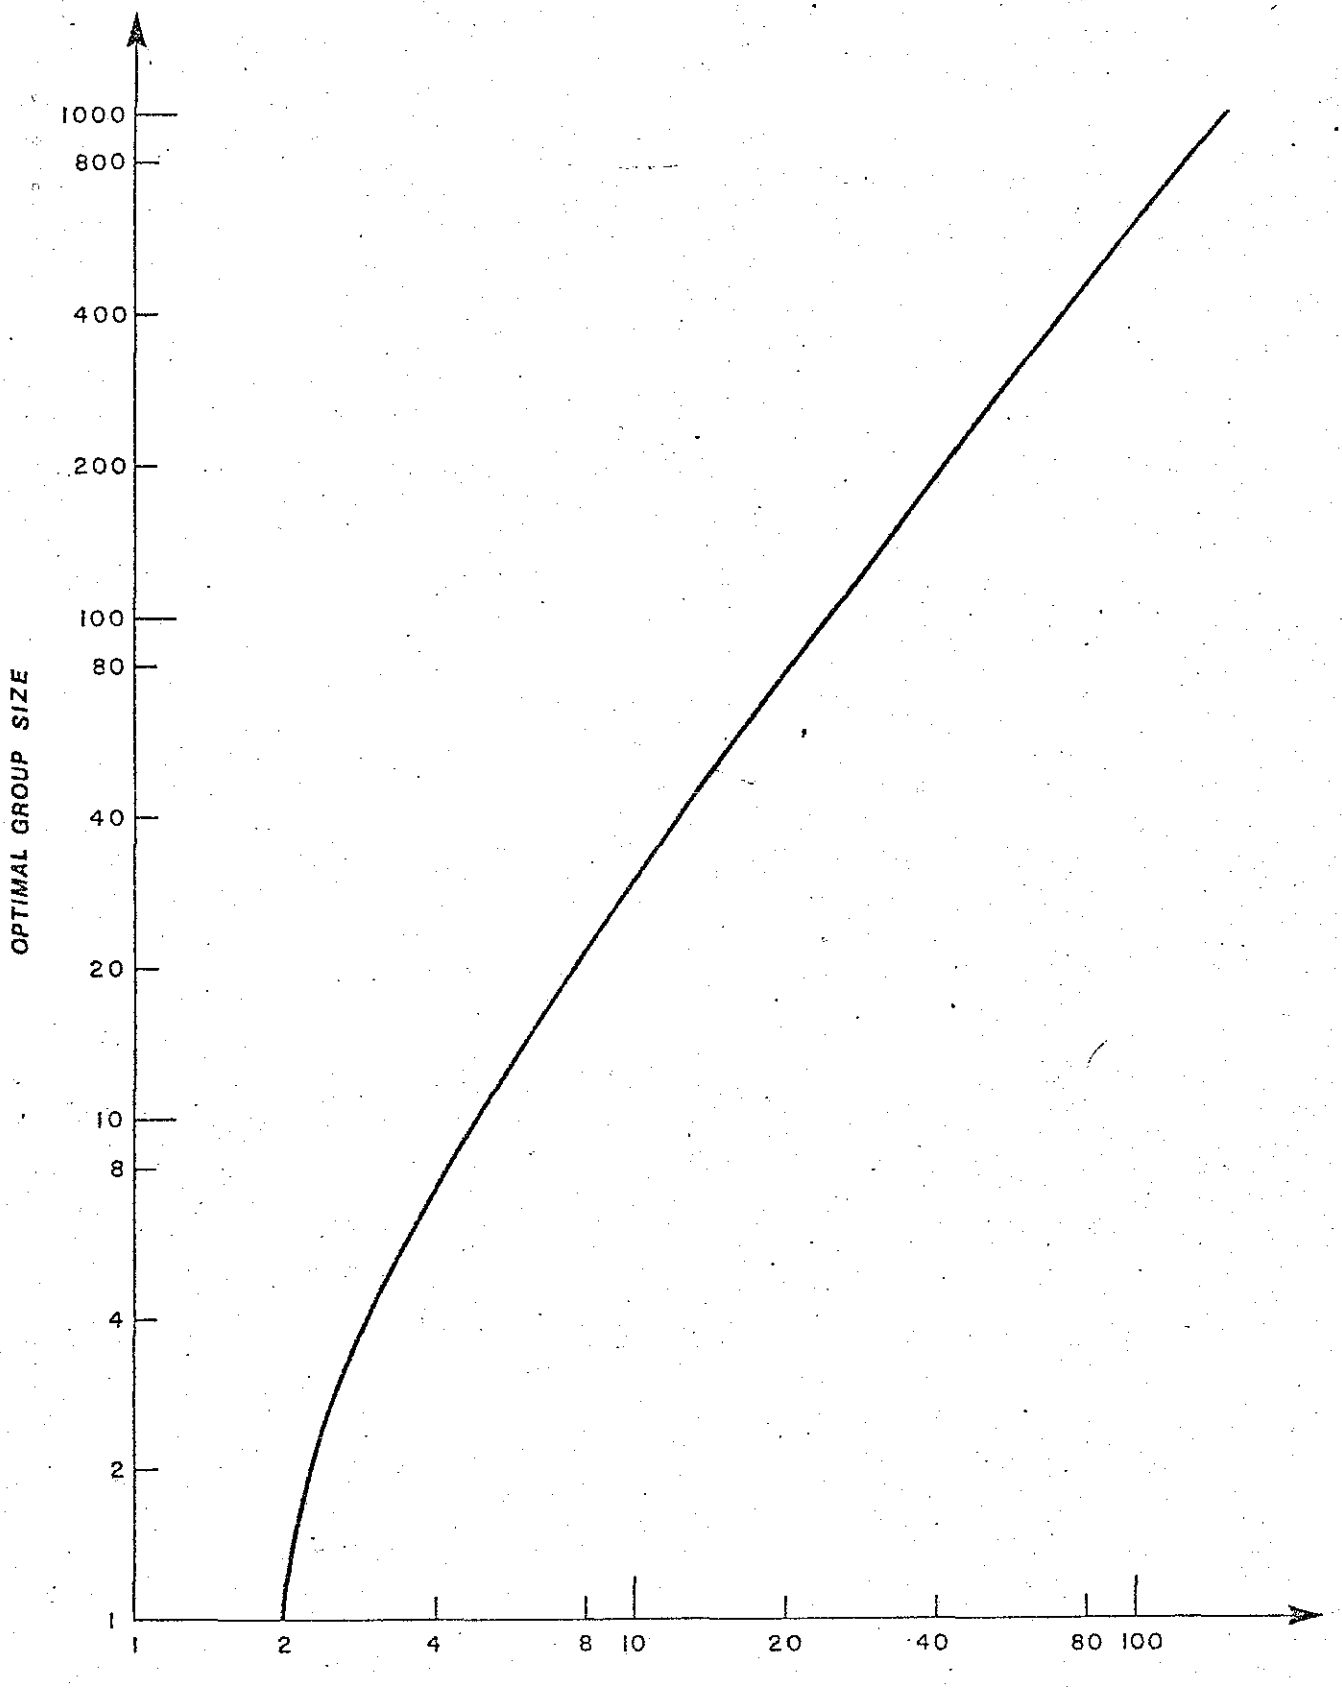


Figure 3. Optimal class size.

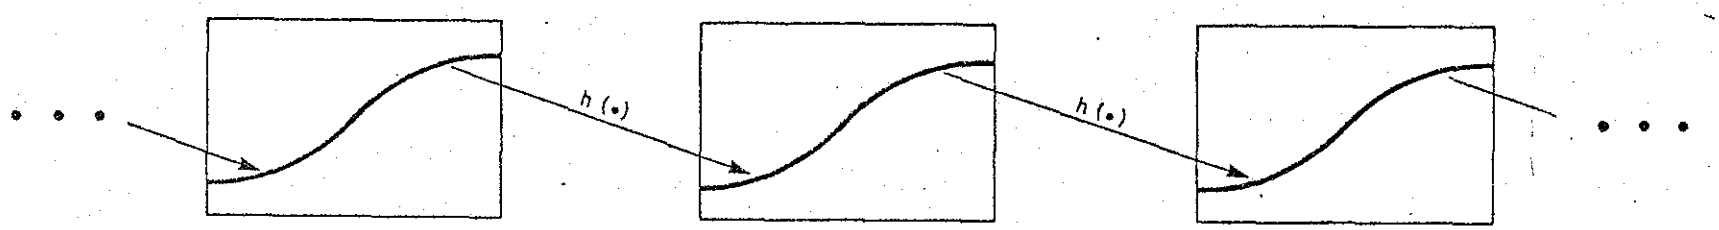


Figure 4. Infinite sequence of similar blocks.

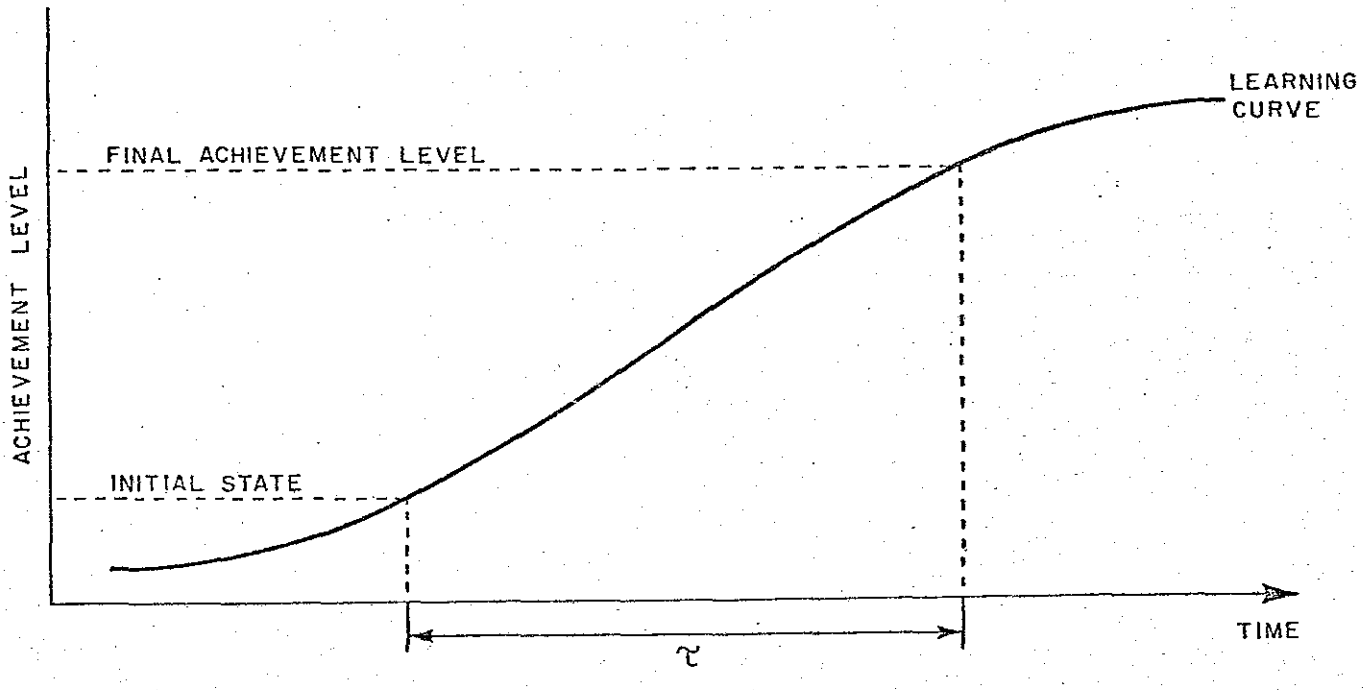


Figure 5. Illustration of steady state condition.

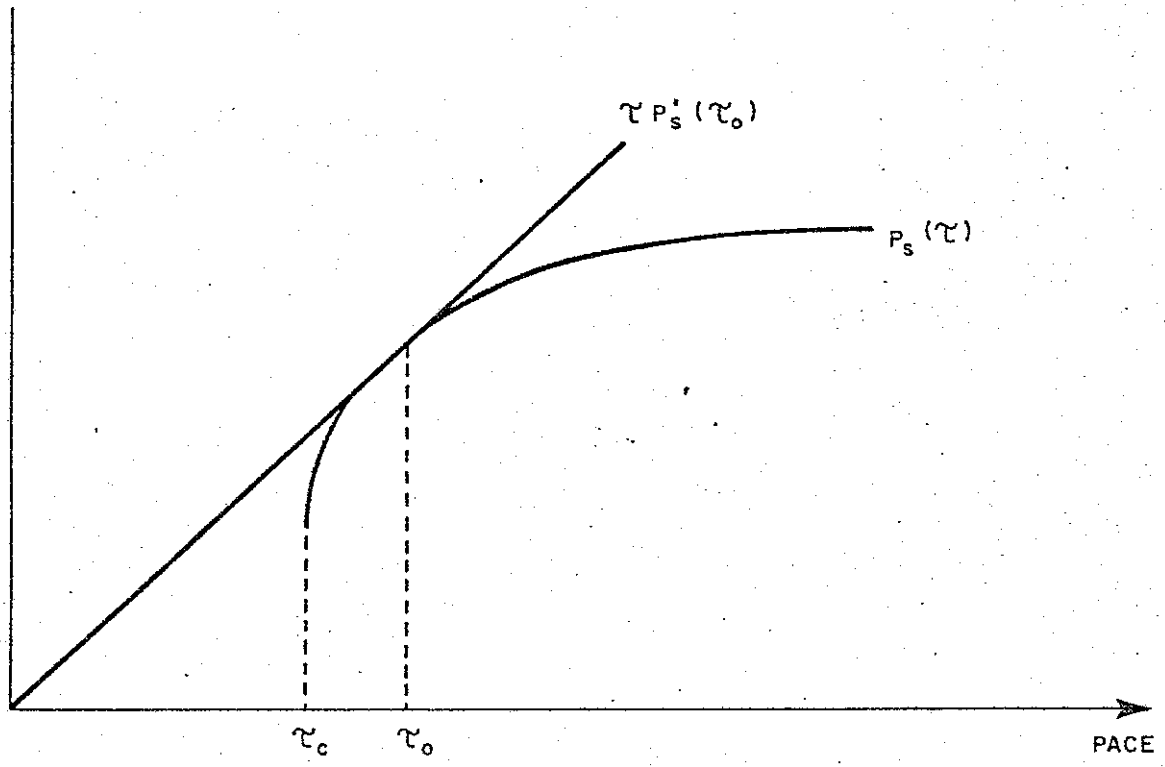


Figure 6. Illustration of optimal pace condition for individual learner.