R E P O R T   R E S U M E S

ED 010 518   24
VARIABLES IN "DISCOVERY LEARNING."
BY- GLASER, ROBERT
PITTSBURGH UNIV., PA., LEARNING RES. AND DEV. CTR.
REPORT NUMBER BR-5-0253-REPRINT-8   PUB DATE   66
EDRS PRICE MF-$0.09 HC-$1.04   26P.

DESCRIPTORS- *DISCOVERY PROCESSES, LEARNING ACTIVITIES, *LEARNING PROCESSES, LEARNING THEORIES, CURIOSITY, MOTIVATION, *BEHAVIOR PATTERNS, LOGICAL THINKING, COGNITIVE PROCESSES, LOGIC, *INDUCTIVE METHODS, *TEACHING METHODS, *INSTRUCTIONAL DESIGN, PITTSBURGH, PENNSYLVANIA, DISCOVERY LEARNING

A PRESENTATION WAS MADE OF THE ANALYSIS OF BEHAVIOR THAT IS REQUIRED AS A FIRST STEP IN THE PROCESS OF DEVELOPING PROCEDURES AND MATERIALS FOR "DISCOVERY LEARNING." TEACHING BY THE DISCOVERY METHOD IS DESCRIBED AS REQUIRING THAT A MINIMUM OF STRUCTURED INSTRUCTIONAL SEQUENCE BE IMPOSED TO ALLOW THE CHILD TO (1) LEARN BY DISCOVERY AND (2) LEARN TO DISCOVER. THE KINDS OF SEQUENCES THAT COULD BE USED WERE GROUPED AS (1) INDUCTIVE SEQUENCES IN WHICH A RULE WOULD BE PRESENTED BEFORE THE EXAMPLES, AND (2) TRIAL-AND-ERROR LEARNING SEQUENCES WHICH WOULD REQUIRE THAT STUDENTS DISCOVER THINGS FOR THEMSELVES. THE INCOMPATIBILITY BETWEEN THIS SELF-DISCOVERY PROCESS AND ERROR MINIMIZATION WAS ALSO DISCUSSED. CONCLUSIONS WERE (1) "LEARNING-BY-DISCOVERY" APPEARS TO INVOLVE NOT ONLY THE PROPERTIES OF INDUCTION AND ERRORFUL LEARNING, BUT ALSO INTERACTIVE EFFECTS WITH TASK PROPERTIES, AND (2) INTERESTING LEADS FOR STUDYING "LEARNING-TO-DISCOVER" MAY BE FOUND IN OPERANT ANALYSIS, COGNITIVE SIMULATION, AND STUDIES OF EXPLORATORY DRIVE. (AL)

5-0253
Reprint (8)

UNIVERSITY OF PITTSBURGH - LEARNING R & D CENTER

REPRINT 8

VARIABLES IN "DISCOVERY LEARNING"

ROBERT GLASER

VARIABLES IN "DISCOVERY LEARNING"

Robert Glaser
Learning Research and Development Center
University of Pittsburgh

[Preprint of a chapter to appear in L. Shulman and E. R. Keislar (Eds.) Learning by Discovery. New York: Rand-McNally, 1966. (in press)]

The research reported herein was performed pursuant to a contract with the Office of Education, U. S. Department of Health, Education, and Welfare. Contractors undertaking such projects under Government sponsorship are encouraged to express freely their professional judgment in the conduct of the project. Points of view or opinions do not, therefore, necessarily represent official Office of Education position or policy.

VARIABLES IN "DISCOVERY LEARNING"[1]

Robert Glaser
University of Pittsburgh

The question assigned to me to consider in this paper asks the following: "Is learning by discovery an important principle in curriculum development?" I approach this task as someone concerned with the design of educational practice. I am interested in the requirements and specifications for the development of procedures and materials for "discovery learning." As an educational designer, I work as a technologist, supplied with a presently meager, but apparently increasing, body of technological principles and practices. These are emerging from the interplay between practical attempts at education and relevant research and knowledge from the sciences which contribute to pedagogical methods.

My design orientation provides me with the following plan of operation: First, I must analyze the behavior with which I am concerned and specify some performance which will represent a standard of competence to be attained at the end of a sequence of educational experiences. This performance specification establishes a model or standard around which individual differences will be displayed. The selected performance must be specified in terms of its class properties because the stimulus, response, and structural characteristics of the subject matter content and the behavioral repertoires involved, will determine what I wish to teach and, correspondingly, how it is to be taught. I should not

---

[1] A paper delivered at a conference on "Learning by Discovery," SSRC, New York, January 28-29, 1965.

be too rigid, however, in sticking to an early specification of this performance because certainly the selection of my instructional goals will be influenced by my analysis of the behavior under consideration.

Second, I need to specify the characteristics of the students that I am to teach. These characteristics need to be determined either prior to instruction or in the process of early learning. I shall need to know the extent to which the student has already acquired some of the things to be learned, the extent to which he has the prerequisites for learning the next instructional steps, the extent to which antecedent learning facilitates or interferes with new learning under the conditions I have in mind, and the extent to which an individual can make the necessary sensory discriminations and exhibit motor skills required for initial learning steps.

With information about both the target performance to be attained and the existing pre-instructional behavior, I can proceed from one state to the other. This sets up my third task, which is to guide or allow the student to go from one state of development to another, and I must construct the procedures and materials that I wish to employ in this educational process. As part of this process, I must make provisions for motivational effects, by which I mean providing conditions which will result in the maintenance and extension of the competence being taught.

Finally, I must make provision for assessing and evaluating the nature of the competence achieved by the learner in relation to the performance criteria that have been established.

If this description of the educational process sounds harshly technological, perhaps some elegance has been lost in analysis. But presumably, once the basic techniques are designed, it is time for the practitioner to apply all the artistry and sensitivity he can muster.

In this paper I will consider only the <u>first step</u> of the <u>first task</u>, i.e., the general examination of the behavior under concern prior to an experimental analysis.

I cannot emphasize enough the importance of this first task--the analysis of behavior. I believe that it has been neglected in psychological research and I also believe that it has been the most important element in recent improvements in instruction. In the design of educational programs, analyses of the terminal objectives to be achieved have been a more influential endeavor than manipulations in methods of teaching these objectives. This is probably so because it is the first step in the sequence of tasks in instructional design. (The fascination of Piaget and Geneva School lies, to some extent, in their keen analyses of children's behavior; but they stop short of the succeeding steps in the operational plan for instruction.)

My analysis begins with an examination of the tasks that have been labeled "discovery learning." I find here that I am confronted by a confusion between two different kinds of events. One has to do with learning <u>by</u> discovery, that is, teaching certain objectives by a discovery method; the other has to do with learning <u>to</u> discover, or teaching for a terminal objective which is manifested by the ability to make discoveries.

The most prevalent case, learning by discovery, is defined usually as teaching an association, a concept, or rule which involves "discovery" of the association, concept, or rule. This is contrasted with a more direct instructional sequence in which a discovery method is not employed. And there are variations between these two. When one examines the task situations and instructional sequences that have been called "discovery" and those that have been contrasted with discovery, what are the outstanding features? Two differences are apparent: First, a learning-by-discovery sequence involves induction. This is the procedure of giving exemplars of a more general case which permits the student to induce the general proposition involved. Assessment of attainment is accomplished by testing whether the student has indeed induced the general proposition by getting him to verbalize it, getting him to apply it to certain exemplars in a way that indicates that he knows the general proposition, or by getting the student to generate additional exemplars. Finding the structure in a body of subject matter instances is an example of induction, and the structure eventually discovered is a general proposition characterizing or summarizing the properties of these instances.

Second, in using the discovery method, the imposition of a structured instructional sequence is minimized in order to provide a relatively unguided sequence onto which the individual imposes his own structure. This kind of sequence, of necessity, allows the student to pursue blind alleys and find negative instances; and consequently, he makes some wrong moves or incorrect responses in the process of learning. "Discovering" implies a low probability of making a successful response. Such being the case, errors have a high probability of occurrence.

Discovery sequences can generally be characterized by these two properties: one, inductive sequences, and two, trial and error, or errorful, learning in various degrees.

We should then examine these two processes, induction and errorful learning. Depending upon the behavioral objectives these processes are to teach, that is, whether they are to result in the establishment of associations, a concept, a rule or generalization, these processes can be considered in different ways and can have different merits. However, before considering them with respect to particular terminal behaviors, it is of some use to discuss them generally.

Induction. I begin with the contrast to induction first. This is a teaching sequence in which a rule is presented before exemplars or instances of the rule. This is expository teaching, and in early work with programmed sequences, a rule-example-incomplete-example sequence, appeared to be an excellent method for the efficient introduction of a new rule (Evans, Homme, and Glaser, 1962). The rule-example-incomplete-example presentation has the student working on the example of a new rule very early in his exposure to it. In this sort of sequence, the student is given an explicitly-stated rule and one or more carefully chosen examples before calling for a response by means of an incomplete example. An effective prompt is then set up which minimizes incorrect responses and which provides the student with the reinforcing activity of directly using the rule. Implied here is the rationale that rather than run the risk (at least in the fixed sequence of early program formats) of having the student induce an incorrect rule, it is preferable to state the rule for him explicitly. This philosophy leads to the rejection of inductive presentation. With a rule-example sequence, the student can recognize and apply a rule with proficiency, and often it seems hazardous and slow to approach a rule through induction or through incidental learning.

With rule and then example, the student adopts the expert's carefully chosen statement of a rule rather than using his own more fallible induction-derived statement. The limited range of exemplars in most teaching and textbook situations may make it possible for the student to induce what is essentially an incorrect rule but one which happens to fit all the examples presented. This is another possible source of danger in the induction process.

The rule-example expository sequence just described is very frequently used in education. A teacher will typically enunciate a principle and follow this with a series of instances of the principle. This is a prevalent procedure because it leads to quick reinforcement for the teacher and the student. They both see close-to-criterion behavior occur rapidly. It is reinforcing perhaps for the same reason that the use of punishment is reinforcing to the teacher--because it brings quick results. Other means of influencing behavior are more laborious and their results only show up in the long run.

Presenting rules first is also very effective because it is more useful to remember a general statement that "mammals are warm-blooded animals and bear their young live" than it is to remember that each specific species, such as monkeys, horses, cows, cats, dogs, etc. is a mammal (Mechner, 1961). Similarly, it is more useful to remember that the square of any number ending in 5 is equal to "x times (x + 1) followed by 25," than it is to remember the squares of specific numbers. In general, one is better off remembering information to be used when it is stored in condensed abstract form rather than in many specific instances. The general statement is often the first one given in teaching because it is easiest to remember and because defining and presenting an adequate sample of instances is a difficult task. Sometimes examples of a general case have little dignity and statement of the rule is more profound. It is often easier, more dignified, and provides more instant knowledge to state the rule before giving examples. While the words I use here suggest negative emotional loadings, nevertheless, for some purposes and for teaching certain kinds of tasks, rule-example is quite effective.

Consider now inductive teaching. Francis Mechner (1961) lucidly points out that great teachers and great writers know the principles of inductive teaching intuitively. Their writings provide us with demonstrations of the effectiveness of giving examples before rules. LaFontaine teaches a code of ethics through a series of allegorical fables. Shaw, in his The Black Girl in Search of God, makes a general point by providing a succession of specific instances which permit the reader to induce the general concept. Interestingly, Shaw's episodes describe non-instances of the concept being established: Moses, Freud, Pavlov, and others are instances of what God is not, and through these examples Shaw conveys his message.

Good writers ingeniously use a series of incidents to establish the concept of a character. It is hard work, at least for me, to read through the development of a Dostoevsky character so that the concept of this character emerges. C. P. Snow, with an inferior literary style, reinforces me more quickly because he tells me that Arthur Brown is this kind of a character--I get the rule first. Induction is also used by poets and composers when they develop general concepts by specific examples of images and themes.

In summary, a long-standing procedure, recognized in society for its excellence, is that concepts and principles are learned by the presentation of specific instances which permit the learner or the recipient to generalize among specific instances of a class and discriminate between instances and non-instances of a class. In these sequences of induction, the learner makes some false inductions, errors if you will, in the course of inducing the rule. Depending upon the subject matter, some rules can be pretty definitively learned, that is, subject to little further correction, such as inducing the concept of equality in mathematics. Other inductive sequences are subject to constant emendation or revision, such as the personality of a character in a story or such scientific concepts as force, energy or the electron.

This inductive procedure is somewhat similar to the way we teach a concept according to the notions of Hull, Skinner, Keller and Schoenfeld, and, I suspect, Piaget. To teach a child the concept of redness, we first insure that the child has a relevant response available, in this case, that he can already say the word "red." However, he does not yet use it appropriately. (It is not under appropriate "stimulus control.") The teaching sequence might point to a succession of pictures or objects asking each time, "What color is this?" Every time the child gives the right answer, he is provided with some event or context which provides confirmation or other reinforcement. The teacher or teaching sequence does not give the rule by pointing to objects and saying this is red, this is green, etc. The child is permitted to make responses by himself to the separate instances. The teaching sequence utilizes various kinds of red objects so that the student is provided with a succession of situations in which a correct response has a high probability of occurrence. Sequences of non-instances or negative instances are employed in which non-red (or possibly the color of the non-red object) is accepted as an appropriate response. The teaching procedure is careful to randomize the non-relevant dimensions involved so that there is included large and small objects, distant and near objects, dark and light objects, and coarse and smooth objects. The child thus learns to generalize among objects in the class of redness. In the course of this process the teaching sequence might introduce obviously different colors asking which one is blue and so on. Once the child says red only to red objects and not to non-red objects, and blue only to blue objects, he has acquired and perhaps discovered the concepts of redness and blueness.

In summary, then, the principle is that an abstract or general case is learned by the establishment of generalizations among specific instances of a class and the discriminations between instances and non-instances of a class. In learning the concepts of triangle and quadrilateral, the student must generalize the response "triangle" to any three-sided figure and the response "quadrilateral" to any four-sided figure. He must also learn to discriminate between these two classes. In larger sequences of topics, a student learns what an operation is after he can add, subtract, multiply, etc. He learns what a proof is after he has seen a large number of different kinds of proofs. And he understands what homeostasis is after he learns about different kinds of physiological equilibria. General understanding is induced from a wealth of experience with specific cases (vide Mechner, 1961).

<u>Errorful learning</u>. The second identifying characteristic of discovery learning is that in the course of discovering things for themselves, students will undoubtedly make mistakes as a result of exploring blind alleys and negative instances. Since it seems that the most intellectually satisfying discoveries are those which are not obvious from the data at the student's immediate disposal, there is the probability that such discoveries will not be made. To this extent, there may be a basic incompatibility between inducing discoveries and minimizing error.

To begin with a contrast again, the development of teaching machines has emphasized the minimization of errors. And while, so far as I know, completely errorless learning has not been demonstrated in a teaching machine program, it has been demonstrated in the ingenious work reported by Terrace (1963a; 1963b) in teaching pigeons a red-green discrimination, and also to discriminate between a horizontal and a vertical line. An error is defined in this work as a response to a stimulus correlated with nonreinforcement, a so-called S-. The results of these studies indicate that performance following discrimination learning without errors lacks three characteristics that were found following learning with errors. Only those birds that learned the discrimination with errors show (1) "emotional" responses in the presence of S-, (2) occasional bursts of responses to the incorrect stimulus, and (3) less effective transfer to related discriminations. The technique Terrace used was to begin with two stimuli, widely separated on three stimulus dimensions, and then to progressively reduce the differences between two dimensions, maintaining only the difference in the third. This technique was recognized by William James in a discussion of discrimination in psychophysics in 1890.

In 1943, Schlosberg and Solomon (1943) reported a study in which they trained rats on a Lashley jumping stand in a black-white discrimination. In order to equate for what they call the "negative factor" which prevents learning and increases response latency, they permitted no errors to occur which would be punished. In this way the value of the negative factor would be determined only by the distance to be jumped and hence equal for all stimulus presentations. They trained their animals very gradually in positive responses to a white stimulus so that an error was never made. As a result of this procedure, the experimenters say that "the gradient established by reward was uncomplicated by the effects of 'punishment'" (p. 26).

As Terrace points out, the demonstration of errorless learning suggests possible revision in currently accepted accounts of discrimination learning. These currently accepted accounts agree that the extinction of responding to S-, and hence the occurrence of errors, is a necessary condition for formation of a discrimination. As succinctly stated by Keller and Schoenfeld, "Extinction is the hallmark of discrimination" (1950). The accounts of Spence and Hull on discrimination learning are similarly based upon learning that occurs in the presence of S and extinction in the presence of S- respectively. Harlow (1959) expressly incorporates error in his error factor theory. In general, discrimination learning without errors is excluded from these conditioning-extinction theories where excitatory and inhibitory gradients are postulated.

The general rationale for error minimization in instruction is the following: (1) When errors occur, there is lack of control over the learning process and opportunity is provided for the intermittent reinforcement of incorrect responses; this results in interference effects highly resistant to extinction. (2) Frustration and emotional effects, that are difficult to control, are associated with extinction and interference. And (3) richer learning, that is, richer in associations, takes place when the associative history of the learner is employed to extend his learning; this is accomplished by mediators or thematic promptings which make positive use of existing knowledge and serve to guide learning.

There is perhaps another reason behind the drive to minimize errors. This is the fact that the use of errors and the possible value of incorrect responses has not been investigated much in studies of learning related to the educational process. The contingencies generally studied have been those following correct responding--a reinforcing event, a punishing event, or withholding a reinforcing event. The contingencies following an incorrect response that have been studied primarily are punishment, withholding reinforcement, and to some extent, variations of corrective feedback information. This latter contingency has not, however, been as systematically investigated as the others. Recent experimental studies like those of Suppes and Ginsberg (1962) suggest that overt correction of errors in young children results in faster learning than does just knowledge about whether or not responses were correct. Although with adults, Suppes points out, studies like that of Burke, Estes, and Hellyer (1954) show that requiring the learner to make an overt correction response after informational feedback does not increase learning rate nor asymptotic performance. (Indeed, many learning studies assume, e.g., Bower (1962), that under certain conditions, correction following an incorrect response has reinforcing value equal to confirmation following a correct response.) In general, the "guided" aspects of studies of guided discovery attempt to make use of error, but effective use requires development of theory and data about the function of error responses.

An exception to the lack of use of error responses in the course of instruction has been the work reported by Lewis and Pask on adaptive teaching systems (1965). The adaptive teaching procedure which these men propose requires the student to reveal, by making some sort of error, the kind of instruction he should receive next. This requirement, they suggest, need not conjure up an image of an aversive and threatening situation in which the student is forced to reveal his ignorance. If adaptive control is competently designed, student weaknesses are revealed by his selection of response alternatives. Where no adaptive procedures are available for dealing with error, the minimization of error is forced upon a teaching procedure.

Error minimization advocates might suggest that the adaptive system could do better by preventing errors from occurring in the first place. Lewis and Pask react to this by pointing out that the presence of error is tacitly acknowledged by the error minimizers when they cue or prompt in the course of a program to adjust a program to the population of students being taught. These non-adaptive programs remove error factors without allowing them to be manifested in the form of overt mistakes. This necessarily involves working in the dark, and hence programs which forestall error often make provisions for far more error possibilities than any one student is likely to have, and, hence consist of less-than-challenging tasks.

Adaptive teaching systems, in contrast to error minimization, take seriously the view that students profit from their mistakes. In addition, an instructional sequence should require that students discover things for themselves, and in the course of self-discovery the student will undoubtedly make mistakes. Thus, there is a basic incompatibility between this self-discovery process and error minimization.

At this point, I am sure of one thing: that is that I have not resolved any issues. But I do hope that I have laid out for inspection what seem to me to be relevant variables and nuances that are involved in "learning by discovery." I have said that the hallmarks in this kind of learning, from my review of the kind of learning situations that have been included under this label, involve two identifying characteristics: induction and errorful learning. And I have attempted to look further into these two aspects to provide some specific variables for their operational handling.

So far I have discussed the characteristics of learning by discovery _in general_. This really does not get us very far in efforts at instructional design because the characteristics of a teaching sequence interact with the properties of the terminal tasks that are being taught. Therefore, induction and errorful learning take on differential usefulness depending upon whether we are teaching response precision, simple associations, concepts, rules and principles, or higher-order strategies.

Consider first the establishment of response precision. An evident characteristic of the educational process that leads to subject matter mastery is the increasing precision of the student's response. The student's initial performance is variable, crude, and rarely meets the criteria of subject matter competence. An effective instructional procedure tolerates this initial state and gradually takes him toward mastery. In order to accomplish this, the teaching process involves the progressive establishment of narrower limits for correct performance. Increasing competence in performing such new skills as learning to write, or learning precise timing in music, is accomplished by gradually contracting performance tolerances. This can be done progressively so that each successive range of successful performance includes a major portion of the range of variations already in the student's performance. Over the sequence of instruction, the range of observed performance will align itself with a particular range of acceptable performance defined as subject matter competence. I would suggest that this can be done and should be done with a minimum amount of errors, since a sudden or inappropriate constriction of performance criteria can lead to extinction and loss of motivation. The use of errors and induction for this kind of learning seems not especially appropriate.

Consider next the learning of associations. The process involved is attaching the increasingly precise responses being learned to particular subject matter stimulus situations so that subject matter mastery is attained because the precise responses are under the command of detailed subject matter discriminative stimuli. In learning translation in a second language, for example, the precise expression of a word meaning already in the learner's repertoire is transferred to new subject matter stimuli. The transfer of stimulus control is a major process in teaching students to make precise subject matter discriminations and teaching them to use previously learned skills in response to new stimuli. This process, like the establishment of response precision, does not seem to require induction and errorful learning. Through the use of mediators, associations can be taught so that errors are minimized. As Gilbert (1960) has pointed out, in learning the correspondence between "one" and the color brown, and "zero" and the color black, in teaching the resistor color code, it is possible to teach such associations on almost one trial by the use of mediating stimuli. The student learns by means of statements such as "one brown penny" and "zero black nothingness." The procedure of stimulus fading used widely in operant conditioning and in programmed instruction can also be used to transfer existing associations to new ones with little error.

With respect to the learning of concepts, I have already indicated that induction may be a useful procedure. Whether or not errorful learning is useful may be debated in light of the work of Terrace in discrimination learning. The question here centers around whether, in the discrimination training aspects of concept formation, errorful learning has useful consequences.

In the learning of rules, principles, and higher order strategies, I can be less than definitive although the use induction and error seem indicated. At this point it is a matter of some systematization and experimental analysis along the lines of the variables I have been discussing.

In summary, learning by discovery, when analyzed, appears to involve not only the properties of induction and errorful learning, but also interactive effects with task properties.

I turn now, briefly, to teaching for discovery, or learning _to_ discover. This is much easier to talk about. Since we know so little about it, one can say anything and enjoy his own speculations without the constraints of knowledge. Again the problem is analysis or the behavior involved because it is unlikely that we can teach discovery behavior adequately until the component behavior repertoires have been analyzed. Once we specify or at least postulate relevant situational and behavioral variables, discovery behavior should be more amenable to instructional manipulation. Right off, it seems likely that discovery behavior is specific to the subject matter domain in which discovery takes place. Discovery requires different concepts and logical combinations when one is working in microbiology or botany, or breaking hieroglyphic codes. Presumably, there is some communality of behavioral repertoires, but there will be much significant specific variance.

Mechner in a recent chapter on "Science Education and Behavioral Technology" (1965) lists subdivisions of scientific method and research skills which seem to be manageable pieces of discovery in science. He lists such things as deductive reasoning skills, inferential reasoning skills, skill in generating hypotheses, skill in selecting "fruitful" hypotheses, skill in testing hypotheses and deciding which experiments to perform, skill in formulating problems that can be solved by the scientific method, and generalized traits such as patience, perseverance, and curiosity. Skill in generating hypotheses is described as a type of scanning, like free association, except that it involves statements about the world. Each statement is checked against experience until it is refuted. If it is not refuted it becomes a scientific contribution. Mechner writes the following: "The terminal behavior test for this skill would require the student to generate and test, at a certain minimal rate, hypotheses about a universe with which he has had some previous experience. The behavioral technologist, in developing the terminal behavior specifications, would have to (a) make the behavior overt rather than covert, (b) circumscribe the universe for the hypotheses, and (c) circumscribe the range of data against which the successive hypotheses are tested by the student. Here, the use of a computer suggests itself. It should be possible to develop a program for the computer that would make the computer behave like a small, artificial, circumscribed universe. This universe would be described by a set of specific input-output relationships, some of them determined and others probabilistic. The student of "creativity" [discovery] would start out by learning his "subject matter," i.e., how this computer behaves under various specific conditions. At the end of this subject-matter training, there would still be a great deal about the computer-universe that would be unknown to him. Here he must begin to generate hypotheses and test them. A program could be designed to develop this skill. The program would teach the student in the standard step-by-step fashion how to generate hypotheses on the basis of some available data and then to test the hypotheses in brute force manner against other available data until the hypothesis is refuted or until the data are exhausted. At the end of the program, the student would be generating and testing hypotheses at the desired rate."

Another way of looking at discovery behavior is suggested by the work on computer simulation of information processing. The work pioneered by Simon and Newell (1964) identifies elementary information processes which combine into compound processes. The processes identified might be considered as heuristics which are teachable behaviors which can be combined by the learner to produce discovery behavior. For example, a paper by Simon and Simon (1962) identifies certain heuristics for discovering and verifying mating combinations in chess. The discovery heuristic involves a tree of move possibilities which explores branches that turn out to be false leads. This "exploration tree" consists of move possibilities in which the attacker has to discover a branching sequence of moves, one sub-tree of which leads to a checkmate. The authors write that, "The exploration tree is precisely analogous to the paths tried out by a subject in a maze-running experiment, except that it includes branches for defender's choices as well as branches for the attacker's trees" (p. 427). A heuristic program simple enough to be simulated by hand is able to discover mating combinations in about 52 of the 136 chess positions. Slight modification of the program adds 10 more mating combinations that would be discovered. Simon and Simon conclude: "The conclusion we reach from our investigations is that the discovery of 'deep' mating combinations by expert chess players requires neither prodigious memory, ultra-rapid processing capacities, nor flashes of insight. Combinations as difficult as any that have been recorded in chess history will be discovered by the selective heuristics we have outlined. . . .

The evidence strongly suggests that expert chess players discover combinations because their programs incorpo[illegible]ate powerful selective heuristics and not because they think faster or memorize better than other people" (1962, p. 429).

A third line of endeavor which is of interest for consideration as an influential variable in teaching for discovery has to do with the study of curiosity and exploration. An increasing amount of research has been directed to the study of this area in the past decade (Fowler, 1965). Research, much of it with infra-human organisms, has indicated that the strength of exploratory behavior is positively related, within limits, to the degree of change in the stimulus situation provided by novel, unfamiliar, complex, surprising, or incongruous situations introduced into the environment. Too great or too abrupt a change, however, is disrupting and may preclude exploration. In complex situations, an individual encounters change by way of his interaction with or manipulations of the elements involved. Such interaction provides the stimulus change which can elicit curiosity and exploratory behavior. Investigations have also demonstrated that behaviors are learned that lead to a change in the stimulus display. Thus, in addition to stimulus change eliciting exploratory behavior, experiments show that organisms will respond in order to secure novel, unfamiliar stimuli. In general, these findings demonstrate that stimulus change or sensory variation may be employed to selectively reinforce behaviors which result in change, and that this variation in the situation will serve concomitantly to elicit exploratory behavior. When stimulus change is used as a reinforcing stimulus, it seems reasonable to hypothesize that learning variables which influence acquisition and extinction will influence the acquisition and extinction of exploratory and curiosity behavior as they do other learned behavior. This suggests that a student's curiosity and explorations which enhance discovery may be elicited and maintained in an instructional environment which provides for appropriate variation in the stimulus characteristics of the subject materials confronting him.

In summary, interesting leads for studying "learning to discover" come from operant analysis, cognitive simulation and studies of exploratory drive. Finally, the excursion that this paper has taken into the intracies of "discovery learning" brings to mind the admonition of Edward L. Thorndike who wrote the following: ". . . if we avoid thought by loose and empty terms, or if we stay lost in wonder at the extraordinary versatility and inventiveness of the higher forms of learning, we shall never understand man's progress or control his education." E. L. Thorndike, 1913.

References

1. Bower, G. H. An association model for response and training variables in paired-associate learning. Psychol. Rev., 1962, 69, 34-53.

2. Burke, C. J., Estes, W. K., & Hellyer, S. Rate of verbal conditionins in relation to stimulus variability. J. exp. Psychol., 1954, 48, 153-161.

3. Evans, J. L., Homme, L. E., & Glaser, R. The ruleg system for the construction of programmed verbal learning sequences. J. educ. Res., 1962, 55, 513-518.

4. Fowler, H. Curiosity and exploratory behavior. New York: Macmillan, 1965. (In press)

5. Gilbert, T. F. An early approximation to principles of progr mming continuous discourse, self-instructional materials. In A. A. Lumsdaine & R. Glaser (Eds.), Teaching machines and programmed learning. Washington: National Education Association, 1960. Pp. 63-634.

6. Harlow, H. F. Learning set and error factor theory. In S. Koch (Ed.), Psychology, a study of a science. Vol. 2. New York: McGraw-Hill, 1959. Pp. 429-537.

7. Keller, F. S. & Schoenfeld, W. N. Principles of psychology. New York: Appleton-Century-Crofts, 1950.

8. Lewis, B. N. & Pask, G. The theory and practice of adaptive teaching systems. in R. Glaser (Ed.), Teaching machines and programed learning, II: data and directions. Washington: National Education Association, 1965. (In press)

9. Mechner, F. Programming for automated instruction. New York: Basic Systems Systems, 1961. (Mimeo.)

10. Mechner, F. Science education and behavioral technlology. In R. Glaser (Ed.), Teaching machines and programed learning, II: data and directions. Washington: National Education Association, 1965. (In press)

11. Schlosberg, H. & Solomon, R. L. Latency of response in a choice discrimination. J. exp. Psychol., 1943, 33, 22-39.

12. Simon, H. A. & Newell, A. Information processing in computer and man. Amer. Scientist, 1964, 52, 281-300.

13. Simon, H. A. & Simon, P. A. Trial and error search in solving difficult problems: evidence from the game of chess. Behavioral Sci., 1962, 7, 425-429.

14. Suppes, P. & Ginsberg, Rose. Application of a stimulus sampling model to children's concept formation with and without overt correction responses. J. exp. Psychol., 1962, 63, 330-336.

15. Terrace, H. S. Discrimination learning with and without "errors." J. exp. Anal. behav., 1963, 6, 1-27. (a)

16. Terrace, H. S. Errorless transfer of a discrimination across two continua. J. exp. Anal. behav., 6, 223-232. (b)