Title

Computer-based Instruction in Initial Reading - A Progress Report on the Stanford Project

Authors

Wilson, H. A.

Atkinson, Richard C.

Publication Date

1967-08-25

Peer reviewed

COMPUTER-BASED INSTRUCTION IN INITIAL READING:

A PROGRESS REPORT ON THE STANFORD PROJECT

by

H. A. Wilson and R. C. Atkinson

TECHNICAL REPORT NO. 119

August 25, 1967

PSYCHOLOGY SERIES

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

COMPUTER-BASED INSTRUCTION IN INITIAL READING: A PROGRESS REPORT ON THE STANFORD PROJECT*

H. A. Wilson and R. C. Atkinson

Stanford University

INTRODUCTION

In 1964 The Institute for Mathematical Studies in the Social Sciences at Stanford University received a grant from the U.S. Office of Education to develop and implement a computer-assisted instruction (CAI) program in initial reading and mathematics. The project, known as the Stanford CAI project, is a continuation of earlier research by Patrick Suppes and Richard Atkinson in the area of CAI; Suppes is responsible for the development of the mathematics curriculum and Atkinson for initial reading. At the beginning of the project two major hurdles had to be overcome. There were no lesson materials in either mathematics or reading suitable for CAI presentation, and an integrated tutorial CAI system had not yet been designed and produced by a single manufacturer. The development of the curricula and the development of the system have been carried on as parallel efforts over the past three years with each having a decided influence on the other.

This paper is a report of the progress of the Stanford CAI Reading Program with particular reference to the school year 1966-67 when the initial classes of first-grade students received a major portion of their daily reading instruction on a CAI tutorial system. This first year's operations must be considered

*This research was supported in part by U.S. Office of Education Grant, OEC-4-6-061493-2089, and U.S. Public Health Computing Grant, USPHS-MH-6154.

essentially as an extended debugging of both the system and the lesson materials. Nevertheless, some interesting comments can be made on the results of this shake-down experience in regard to both the feasibility of CAI instruction for first-grade students and on the impact of such instruction on the overall learning process.

GENERAL DISCUSSION OF CAI

Before describing the Stanford Project, a few general remarks about computer-assisted instruction and CAI systems might be in order. Three levels of computer-assisted instruction may be defined. Discrimination between levels is based not on hardware and software considerations but primarily on the complexity and sophistication of the student-system interaction. The most advanced student-system interaction may be achieved with a simple teletype terminal and the most primitive interaction may require some highly complex programming.

At the basic interaction level are those systems which present a fixed, linear sequence of problems. Student errors may be corrected in a variety of ways (e.g., a prompt may be given in the form of a partial answer, or the entire correct response may be furnished following an error) but no real-time decisions are made for introducing unique teaching strategies or instructional materials on the basis of a student response. Such systems have been termed "drill and practice" systems and are exemplified by the Stanford Drill and Practice Program in Arithmetic and in Spelling.

At the other extreme of the interaction scale are "dialogue" programs of the type under investigation at Bolt, Beranak and Newman and at Stanford University. The goal of the "dialogue" approach is to provide the richest possible student-system interaction where the student is free to construct unrestricted natural-language responses, ask questions and, in general, exercise almost complete control over the sequence of learning events.

"Tutorial" programs lie between the above extremes of student-system interaction. Tutorial programs have the capability for real-time decision-making and instructional branching contingent on a single response or on some subset of a student's response history. Such programs allow the students to follow separate and diverse pathways through the curriculum based on their individual performance patterns. The probability is high in a tutorial program that no two students will encounter exactly the same sequence of lesson materials. However, student responses are somewhat limited since they must be chosen from a prescribed set of responses or constructed in such a manner that relatively simple text analysis will be sufficient for their evaluation. The Stanford CAI Reading Program is tutorial in nature and it is this level of student-system interaction that will be discussed in this paper.

CAI: A Tool for Teaching.

Computer-assisted instruction has a theoretical basis, as indeed does all programmed learning, in the notion that immediate reinforcement facilitates learning. For the human learner, reinforcement may come through verbal praise acting as a reward or through simple knowledge of the results of one's own actions. In the reading program correct responses are rewarded by such verbal messages as "good," "you're doing fine," "right," etc. Since studies have shown that reinforcement tends to lose its effectiveness when it is continuous and repetitious, verbal rewards in the reading program are given on an intermittent basis. Immediate feedback is provided through reward messages, through the presentation of the next problem and also through "wrong answer" messages.

Uninformed criticism of computer-assisted instruction often carries an implied component of fear that the instructional process will somehow become dehumanized, that the students will become little automatons themselves. As it turns out, however, the elimination of the social intercourse aspect of learning through CAI is one of its great strengths. The computer is an eternally patient teacher. The machine never becomes angry or threatening. Those of us who have spent some years teaching in the classroom are well aware of the fact that after repeated student errors it is difficult, if not impossible, to restrain certain voice or facial cues which indicate our displeasure. The messages coming from the machine, however, are completely free of any such threat or anger. The "wrong answer" messages recorded in the quiet of the recording studio can be a continuously neutral "No, this is the right word. Touch it."

We have found that attention of even very young students can be maintained at a high level by appropriate pacing of the material. Their attention is also increased by the use of partitions between the response terminals. The individual student is not distracted by the actions of his classmates, nor is he carried along by their responses. Each individual is independently responsible for interacting with the learning materials.

The basic rationale, however, for the use of a computer-assisted instructional system as a teaching device, resides in the system's potential for individualizing instruction. If there is any one fact which has been thoroughly established in sixty years of intensive investigation in education, it is that a wide range of individual differences will be found in any classroom on any dimension one wishes to examine. CAI offers us a tool for tailoring our instructional procedures to these individual differences.

The current paper and pencil programs can accommodate individual differences in learning time. A bright student who responds rapidly can complete many more frames and cover a greater amount of material in a given time interval than can the slow student who responds quite deliberately. However, the nearly unlimited branching capability inherent in a CAI system not only allows the students to proceed independently in terms of speed, but also permits them to follow essentially different paths through the curriculum. Branching decisions for each student in a CAI system may be contingent upon a single response, his past history of responses, response latency, or some combination of these considerations weighted by previously acquired psychometric data.

CAI: A Tool for Research.

The usefulness of CAI as a tool for research resides in two factors: 1) the control of independent variables and 2) the detailed response data which is recorded by the system.

Educational research conducted outside the laboratory in an actual school situation has long been plagued by the impossibility of controlling many variables inherent in classroom organization and in presentation of materials by the classroom teacher. CAI, in a sense, brings the laboratory into the school. Our own CAI laboratory has achieved a degree of control of environment and presentation that has been heretofore impossible in a classroom setting. The temperature and the lighting in the terminal room are constant. The immediate environment of each student's response terminal is precisely the same as every other student's.

The chairs and the machines are identical for all students. Every picture seen in the projection device, every bit of orthography or other display on the scope, and every audio message which the student hears, have been previously specified and can be as standardized or varied as the experimenter desires. This is not to say, of course, that all sources of variation are controlled. The CAI facility does achieve, however, a degree of control equivalent to that of the psychologist's learning research laboratory. Many problems in learning theory, which have been investigated rigorously only in a laboratory setting, can now be looked at in an on-going school context.

The second capability of CAI which is extremely important for research is the collection of fine-grained response data. For example, in the Stanford-Brentwood CAI laboratory, each response that each student makes is recorded on the data tapes. Each response record includes a complete description of the response in terms of coordinates taken from the face of the scope or the keys depressed on the typewriter. The response is defined as correct or incorrect; and if it is incorrect, it is categorized according to the type of error made. The response latency is recorded in tenths of a second. The contents of 31 counters and 32 switches associated with the student's past history of performance are also recorded with each response on the data tape.

Another important use of CAI as a research tool is found in the area of mathematical psychology, particularly in the area of mathematical learning models. The quantity and nature of the response data which may be gathered in the CAI system allows the mathematical psychologist to test his various models in a situation that is a much closer approximation to actual classroom learning than has existed in the past. Typically, such models have been tested through infra-organism behavior or through such contrived tasks as paired-associate list learning or probability learning. The kind of sequential response data which we are gathering in the Brentwood facility will be used in the development of optimization models for learning (Groen and Atkinson (1966); Atkinson and Shiffrin (1967)).

CAI: A Tool for Curriculum Evaluation.

Another feature of CAI in the educational process is in curriculum evaluation. Our current methods of evaluation are extremely gross, relying on standardized tests or specially devised tests given on an intermittent schedule. The best that can be expected from such evaluation procedures is to be able to compare the general outcomes of one method or one curriculum approach to some other method or approach. Little or nothing can be said about the efficiency of any specific section of the curriculum. It is exactly at this detailed level that CAI exhibits its greatest power for evaluation. The performance data gathered in the CAI system may be examined at all levels, from the perspective of overall goals or from the perspective of the various strategies and approaches adopted in the curriculum; we may examine responses to blocks of homogeneous problem types, responses to separate problem types and responses to the individual problems themselves. At each level we can look at the students' performance records to discover if this section or level or item of the curriculum is functioning in the manner for which it was designed.

DESCRIPTION OF SYSTEM

Configuration.

The Stanford-Brentwood laboratory utilizes an IBM 1500 CAI system. The 1500 system was designed and constructed by IBM engineers in close collaboration with Stanford personnel. The student response terminals consist of a cathode ray tube (CRT), a modified typewriter keyboard, a light pen, a film projection device, and a set of earphones with an attached microphone.

The CRT is essentially a television screen on which alpha-numeric characters and a limited set of graphics (i.e., simple line drawings) can be generated under computer control. The film projector is a 16 mm. rear-view filmstrip projector. Still pictures in black and white or color may be displayed under computer control. Each filmstrip, in a self-threading cartridge, contains 1024 frames which may be accessed randomly by means of a binary code along one edge of the film.

The major response device used in the reading program is the light pen. The light pen is a light-sensitive probe which registers a portion of the CRT trace comprising the CRT raster. The precise interval between the initiation of the sweep trace and registry of the trace by the light pen indicates where the light pen has been touched to the screen. This location is stored as a set of coordinates in computer memory. These coordinates are evaluated by the systems program and compared against predefined coordinates in the lesson program. This comparison permits evaluation of a light pen response as correct, incorrect or undefined. Responses may also be entered through the keyboard, however, no use has been made of this response mode in the reading program. This is not to minimize the value of keyboard responses but rather to admit that we have not as yet addressed ourselves to the problem of teaching first-grade children to handle a typewriter keyboard.

Prerecorded audio messages are played to the children through the earphones. There is also an on-line recording capability. The children may, when the microphone is activated, record their own production of a given text displayed on the screen. This can then be played back to them with or without an adult model. The recording and play-back capability helps compensate for the absence of a voice analyzer. The system cannot evaluate the student's vocal production and therefore each student becomes his own voice analyzer.

The 16 student response terminals are serviced by an IBM 1800 Process Control computer. This central processing unit has a relatively limited (i.e., 32K) immediate access core storage. Rapid access bulk storage is provided by six interchangeable disk drives, each disk containing 512,000 sixteen bit words. The audio component of the system consists of a bank of IBM 1505 audio units. Each audio drive unit is connected to one of the response terminals but the connections can be varied at will. Response data flowing into the system from the student terminals is recorded on two IBM 2402 tape units. An IBM 1501 station control, a 1442 card reader punch and a 1443 line printer complete the configuration of the IBM 1500 CAI system.

Insert Figure 1 about here

Time-Sharing.

The IBM 1500 CAI system is a time-sharing system in the sense that the activities of each student response terminal are examined in sequence and appropriate actions taken. The time-sharing is not strictly sequential in that an interrupt feature allows certain priority conditions within the instructional system to alter the sequence of programmed instructions.

A non-technical characterization of the flow of information in the system may serve to suggest the operation of the time-sharing system. Assume that we are breaking into an instructional session and that the system is preparing to examine a response entered by a student at terminal 1. The appropriate terminal record and lesson instructions are read from the disk storage and placed in core. The response coordinates from terminal 1 are evaluated and compared to those stored in the lesson program for the given problem. (See the section on lesson coding for a detailed discussion.) Decisions are made on the basis of the lesson program logic and appropriate commands are given: 1) to the station control to display certain text on the CRT, 2) to the film projector to position and display a certain frame, and 3) to the audio unit to play track α segment n to segment n+p. The attention of the system then moves to terminal 2 and the process is repeated.

The above description is almost allegorical in its oversimplification. It is intended only to give a general feeling for what is meant by time-sharing on the system. For a complete technical discussion the reader is referred to CAI Programming Systems Users Guide, Preliminary Draft, IBM, Armonk, New York, 1967.

The entire process moves at a very rapid rate. Subjectively, the student at the terminal feels that he has the full attention of the system. The response time for the CRT is less than one second and the projector response is nearly as fast. The audio response time is somewhat slower, ranging from 2-4 seconds on the average.

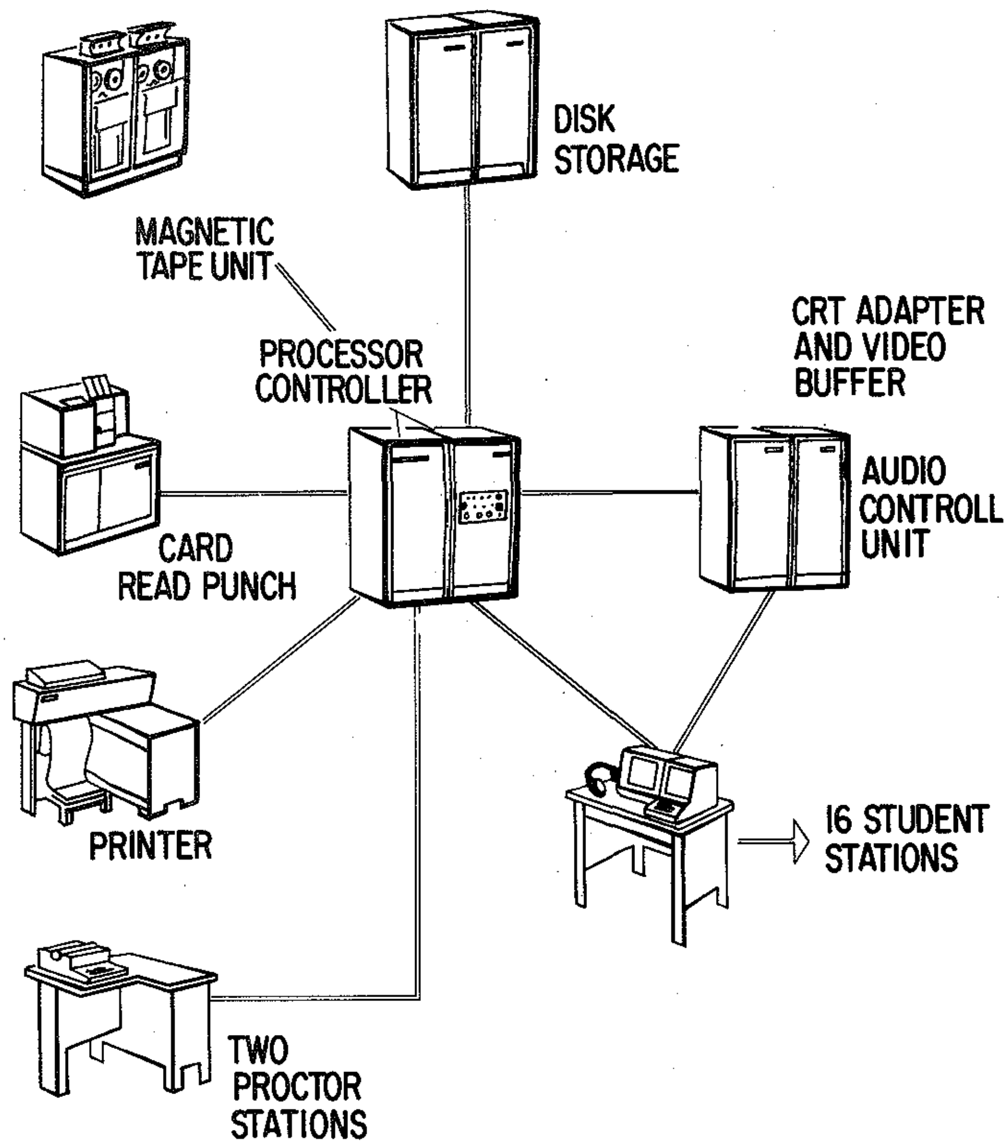

Figure 1.

System Configuration for Stanford-Brentwood CAI Laboratory

SCHOOL-LABORATORY ORGANIZATION

School Population.

The Brentwood School, a K-6 elementary school, is located in the Ravenswood City School District in East Palo Alto, an unincorporated section of San Mateo County adjacent to the City of Palo Alto. The population of the Brentwood School is approximately 80 percent Negro with the remaining 20 percent about evenly divided between Mexican-American, Oriental-American and Caucasian. The school qualifies for federal aid to impoverished areas under Title I.

The Stanford CAI Project is concerned with the first-grade students at Brentwood School. There are four first-grade classrooms at Brentwood, two of which receive instruction in mathematics under computer control and two receive instruction in reading. Of the two classrooms in the reading program, the student population is approximately 90 percent Negro, 6 percent Mexican-American, 2 percent Oriental-American, and 2 percent Caucasian.

Primary-grade Organization.

The students are grouped in first grade at Brentwood according to high/low maturity. The grouping is made on the basis of recommendations by the kindergarten teachers and a reading-readiness test designed by district personnel. The reading project works with both a high maturity and a low maturity room.

Both of the teachers involved in the Stanford-Brentwood CAI Reading Project are experienced first-grade teachers; one has been teaching in the Ravenswood District for 16 years, the other for 7 years.

Preparation for Innovation.

A serious attempt was made on the part of the Stanford Project to prepare the teachers and parents of the Ravenswood District, and of the Brentwood School in particular, for the acceptance of the technological innovation of a computer-assisted instructional laboratory. A good deal of effort was expended during the school year 1965-66 in this preparation. Workshops were held for the Brentwood personnel and an extension course was conducted by the project staff for the district personnel. In both the workshops and the extension course discussions centered on the general concept of programmed learning, linguistics and reading, time-sharing on the computer, a description of the IBM 1500 system, and a very thorough discussion of the reading curriculum. Besides the workshops and the inservice training courses, weekly meetings were held with the principal and primary grade teachers at Brentwood to discuss in greater depth the plans for the ensuing year. During the summer of 1966 a two-week workshop was held on the Stanford campus for the two teachers who would be directly involved in the reading program.

No attempt has been made on the part of the Stanford project to dictate in any way the on-going reading instruction within the classroom. Rather, we have viewed the complete initial reading program as a shared responsibility based on informed cooperation between the project staff and the classroom teachers. Weekly meetings are held between the teachers and the project personnel to evaluate the students' progress on the system and to exchange views and information about both the classroom and the laboratory instruction, and the performance of the students in both environments.

A member of the Stanford staff is permanently placed in the Brentwood School, acting as a school-affairs liaison officer. His duties include the resolution of any organizational problems that arise in the running of the CAI project within the Brentwood School and also the handling of large numbers of visitors who have visited the laboratory during the year.

The preparations and the current efforts for school-project cooperation have paid off handsomely. The enthusiasm and support of the teachers has been highly gratifying. We have also held many exhibits, open houses, and discussion groups for the parents of the Brentwood School and the patrons of the Ravenswood School District. Again, the support of the parents, the school board, the administration and the teachers has exceeded all expectations.

Laboratory Organization.

The laboratory is housed in a rectangular prefabricated steel structure approximately 3200 square feet in area, located on the Brentwood School grounds. The building contains the terminal room, the off-line teaching room, the central computer room, and a group of offices for the laboratory personnel.

The staff of the Stanford-Brentwood CAI laboratory consists of ten members. The laboratory is under the general management of a senior programmer who is also in charge of the data reduction staff. His staff includes two programmers, two graduate students and a secretary who also functions as a receptionist. The systems group is headed by another senior programmer who has on his staff an assistant programmer and a computer operator, plus a technician who handles audio assemblies. The coding groups for both the mathematics and the reading program are also housed in the laboratory. The reading coding group is directed by a senior lesson programmer and consists of four coders, plus part-time debuggers and graduate assistants. A similar number of personnel comprise the mathematics coding group. Three proctors handle the children in the terminal room itself and are responsible for off-line instruction. IBM has also provided several customer engineers who are either on duty at the laboratory or are on call. At present the stability of the system is such that the staff of IBM customer engineers has been reduced to a single man during school hours.

Laboratory Operation.

Full scale operation with the students began on the first of November, 1966. The starting date was somewhat later than originally anticipated, due to a delay in the delivery of the 1500 system which arrived at the Brentwood School on July 10. The ensuing interval from July 10 to November 1 was devoted to a shake-down and debugging of the system which was, of course, untested in on-line operation.

On the first of November the students began coming to the laboratory on a daily basis. The students received a week of orientation lessons outside the terminal room in the off-line teaching room. During this period they were acquainted with the use of the earphones, microphone, and the light pen. Exercises and games were carried out which acquainted them with the equipment and the kinds of learning tasks that they would encounter in the terminal room. During the second week in November they were introduced gradually to the actual response terminals on a staggered basis; that is, a lesson in the terminal room was followed by a lesson in the off-line teaching room. By November 15 the students were on the system daily for a 20-minute instructional period. During these early exposures to the terminal equipment an adult was stationed behind each child to assist him with whatever problems he might find in handling the equipment. It was soon evident that the children were adapting quickly to their new environment, and the adults were gradually withdrawn. By mid-November the proctor staff within the terminal room consisted of the same personnel as it does today -- one teaching proctor, one machine proctor and a remedial reading teacher.

The students come to the laboratory in four groups, since the terminal room is equipped to accommodate only one half of a normal classroom at a time. Each group receives 20 minutes of instruction and there is a minimum of 10 minutes between groups. The 10-minute interval is necessary in order to prepare the system for the next group. The headsets and light pens are swabbed with alcohol to minimize the risk of infection. Each student is signed-on at a response terminal and the appropriate films and audio tapes are loaded into the projectors and audio-drive units. When the sign-on process is complete the name of the student assigned to a given terminal appears on the CRT.

When a new group arrives, the students enter the terminal room, seat themselves before their assigned terminals, put on their earphones, touch their names with the light pen and the lesson begins. During the instructional session the machine proctor is stationed at a proctor's typewriter which transmits messages from the central system. These messages consist primarily of identification of some system or terminal failure, transfer of a student from one lesson to the next, or notification of an error limit exceeded at some terminal. In the case of system or terminal failure the machine proctor takes appropriate action, which may range from transferring a student to an empty terminal to notifying the head systems programmer. When a student transfers to a new lesson the machine proctor must change his audio tape. Each block of problems within a lesson has an error limit which is 50 percent of the total number of problems within the block. If this error limit is exceeded by the student, notification of this condition is transmitted to the machine proctor, who in turn conveys the information to the teaching proctor. The teaching proctor then observes the student and chooses between two courses of action. If the difficulty appears to be mechanical (e.g., the student is not using the light pen correctly) or some minor misunderstanding of directions or lapse of attention, the teaching proctor may sign on the terminal with the student and help him through the troublesome section. By signing on the terminal, a proctor bit is set in the response record to identify responses which are not necessarily those of the student. If, however, the student's difficulty appears to be of a serious nature, the proctor may remove the student from the terminal and transfer him to the remedial teacher for diagnosis and personal instruction in the off-line teaching room. Fortunately, during the entire year's run the necessity for such off-line instruction has been very infrequent. The remedial reading teacher has functioned primarily as an assistant to the teaching proctor within the terminal room and has been brought into play in her major function only when the system has gone down or an individual terminal has failed.

Both the teaching proctor and the remedial reading teacher hold elementary teaching certificates from the State of California and are experienced classroom teachers. The machine proctor is a member of the lesson programming group. The children are accompanied between the classroom and terminal room by an assistant proctor who performs certain clerical tasks while the students are on-line.

Since the two first grades in the reading program are part of the ungraded primary plan at Brentwood, a certain amount of transfer of students between the low and high maturity rooms has taken place in the course of the year. The exchange of the students from one group to another has been one of the major topics of the weekly meetings between the first-grade teachers and the Brentwood project staff. We have been able to accommodate all requests by the teachers for exchanging students within groups. Our only stipulation has been that the exchanges be made known to the project staff by Thursday of the week preceding the actual change of students from one group to another. The lead time has been necessary on our part in order to keep the student registry pack in proper order.

The amount of migration of students to and from the Brentwood School has been relatively small, at least for the two classrooms involved in the CAI reading project. Four of the original students have left the school and have been replaced by four incoming students.

The population of the first-grade classrooms was somewhat smaller than anticipated for this year, which left terminals unoccupied for each of the four groups. During the early phases of the program this arrangement was fortuitous in that spare terminals were available for transferring a student when his original terminal failed. After Easter vacation it was felt that the system had achieved a degree of stability, however, which no longer required a large number of terminals to be available as backups. The empty terminals were then filled with an appropriate number of remedial second-grade students. Since they started late in the year their performance will not be considered in this report except to note in passing that the second-grade teachers involved report a definite increase in interest and application to the reading task within the classroom.

Laboratory-Classroom Cooperation.

It would be unrealistic in the extreme to claim that all initial reading instruction is being carried out in a 20-minute session on the CAI system. As stated previously, we view the program as one of cooperative effort between the normal classroom reading instruction and the Stanford CAI project. Accordingly, we have acquainted the teachers with the CAI curriculum and have issued regular reports on the performance of their students on the system. We have asked, in return, only that the classroom teachers keep us informed of the kind of instruction they are carrying out within their own rooms. This dialogue is effected during our weekly meetings.

Both classroom teachers started the year with an individualized reading program. However, after the Christmas vacation a more structured basal series approach was introduced. From approximately mid-February to the end of the year an eclectic approach was used, incorporating features of both the basal series and the individualized program.

The project's contribution to the dialogue consists primarily of a weekly report which is generated by the computer and which gives the student location by lesson within the curriculum, plus a weighted index of each student's performance within each of the standard lesson blocks. This performance index is cumulative but is weighted for current performance more heavily than for past performance. Also included on the weekly report is a record of time lost from the system due to various causes and a cumulative total of each student's time on the system. On the basis of the weekly report and discussion within the weekly meetings, the classroom teachers are able to devise and carry out instructional procedures appropriate to each student's progress on the system.

Ideally, a CAI system such as this would have available for the classroom teacher a large set of closely correlated materials. Efforts have been initiated toward this goal. An extensive teacher's manual is being written which will include not only a detailed description of the curriculum but also a large number of correlated and well prepared classroom activities from which the teacher may choose. Exercises and games will be keyed to blocks of lessons and will be applicable to small groups as well as to individual students. The manual, at least in its elementary form, will be available for use during the school year 1967-68. Both of the classroom teachers involved in the reading project for this past year will assist in the design and writing of the manual during the summer of 1967.

CURRICULUM RATIONALE

A basic assumption underlying the Stanford CAI Reading Curriculum is that the English speaking child brings to the initial reading task a relatively large vocabulary and at least an operational knowledge of English syntax. He has a knowledge of the language which is sufficient to enable him to communicate with his peers and with adults. Therefore, the primary goal of initial reading is not to teach the language but to teach the orthographic code by which we represent our spoken language.

If one is attempting to teach a code, the most reasonable approach is to begin, not with all the irregularities and exceptions, but rather with the regular and consistent patterns. This position is not unique to the Stanford Project. It has been advocated by linguists for some years (Bloomfield (1942), Fries (1963)) and has been implemented in several "linguistically oriented" reading series. The sequencing of monosyllabic patterns (Table 1) and certain extensions and refinements of the basic notions as stated by the above-mentioned authors, constitute the original work in curriculum design carried out by the Stanford Project.

A detailed discussion of the psycholinguistic rationale of the curriculum may be found in Rodgers (1967) and Hansen and Rodgers (1965). The following quotations from Rodgers (1967) provide a concise summary of our goals and procedures.

> From a practical point of view, our program is an attempt to provide non-readers with some limited analytic skills--phonological, morphological, syntactic and semantic--and some considerable confidence in the use of these skills. It is not our intention to teach the child all of the sound-symbol pattern correspondences, all of the morphological variations, all of the usages of frequent vocabulary items, or all of the sentence patterns of English. It is our intention to give the student enough skill and self confidence to involve him in that confrontation known as beginning reading. We believe it is the ability to make

reasonable inferences concerning unfamiliar or unobserved sequences on the printed page that we are ultimately trying to teach in reading.

We have defined the Stanford approach to initial reading as applied-psycholinguistic. Hypotheses about the nature of the reading process, the nature of learning to read, and the nature of teaching reading have been constructed on the basis of linguistic information about the structure of language, empirical observations of language use, and an analysis of the function of the written code. These hypotheses have then been tested in experimental situations, structured to represent as realistically as possible actual learning and teaching situations. On the basis of experimental findings, these hypotheses have been modified, retested and ultimately incorporated into the curriculum as principles dictating presentation variables and values. This is, of course, somewhat of an idealization since very little curriculum material can be said to have been the perfect end-product of rigorous empirical evaluation. We would claim, however, that the basic tenets of the Stanford program have been formulated and modified on the basis of considerable empirical evidence. It seems probable that these may be further modified or re-formulated on the basis of the considerably greater amount of empirical evidence which will be available as the result of a year's CAI experience with classes of beginning readers.

CURRICULUM DESCRIPTION

The Stanford CAI Reading Curriculum may be divided into four broad areas of concentration: 1) decoding skills, 2) comprehension, 3) games and other motivational devices, and 4) review. The lesson material and teaching strategies will be discussed briefly in each area. While all lessons are not alike in their sequence of events, the block level flow chart (Figure 2) of Level III, Lesson 9, may be considered as a fairly representative example.

Insert Figure 2 about here

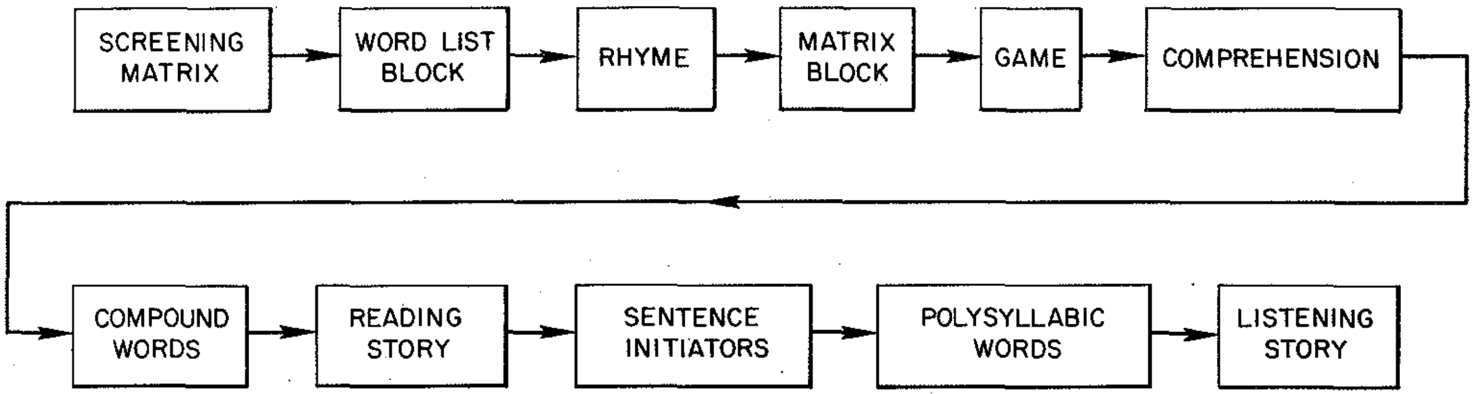

Figure 2.

Lesson Flow-Diagram

1. Decoding Skills

Letter and Letter-String Identification.

No direct attempt is made to teach the names of the letters. From the point of view of our program, letter names are in many instances at odds with the dominant sound which the letter represents. Letter names are used in the audio messages but the emphasis is on discrimination and recognition of individual letters and letter-strings in varying orthographic contexts rather than on mastery of the letter names themselves.

Single Letter Matching. A model letter is presented on the projector and two or three letters are presented on the CRT. The wrong alternatives are designed to include letters considered to be easily confused due to their similar orthographic features. The student is requested to choose the letter on the scope that matches the model presented on the projector.

Letter-String Matching. A model string of two or more letters is presented on the projector and a choice of two or three strings is presented on the CRT. The student is requested to choose the alternative that matches the model string on the projector. The alternative letter-strings include reversals, one letter differences, two letter differences, easily confused letters, and differences in serial order.

Same-Different Task. A pair of letter-strings is presented on the scope. The student is requested to touch one of two boxes that designates whether the letter-strings are the same or different. The letter-strings are designed to include reversals, one letter differences, two letter differences, and easily confused letters. This task provides additional opportunities to acquire functional recognition cues and the serial order concept of letter sequences within words.

Word List Learning.

This section of the curriculum may be described as a set of paired-associate tasks where the stimulus is the verbal pronunciation or pictorial representation of a word (or both), and the response is the correct identification of the appropriate written word in a list of written words. The lists for any given lesson are composed of words generated by the rhyming and alliterative patterns being presented in that lesson.

Quantitative learning models, appropriate to the paradigms in this section but similar in nature to those existing for classic paired-associate learning (Atkinson, Bower, and Crothers, 1965), may be developed to describe acquisition processes and to assess the effects of learning and forgetting.

Five problem types (PT) are included in this section with each problem type containing approximately six problems. Each PT represents a step in a cue-fading technique. The five PT's are:

PT 1 - Cues: Picture, Orthography, and Audio

PT 2 - Cues: Picture and Audio

PT 3 - Cues: Audio only

PT 4 - Cues: Picture only

PT 5 - Criterion Test. Cues: Audio only

The student responds to a set of cues by touching a word in a list of words on the CRT. Each student response receives immediate feedback. If a response is correct it is reinforced. If it is an error the correct answer is indicated by an arrow and an overt correct response is required before the next problem is presented, or the student is branched to appropriate remedial problems. This is

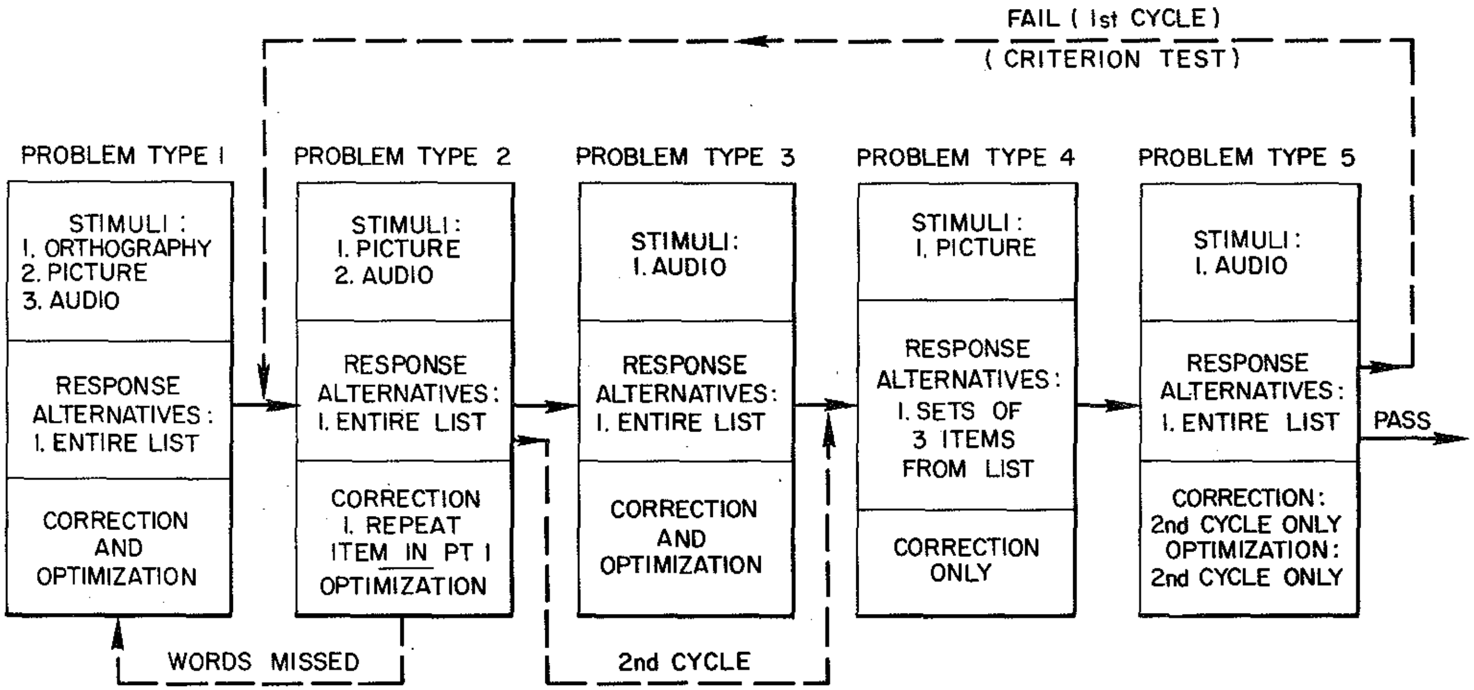

Figure 3.

Flow Diagram for Word List Block

a general teaching strategy which is adhered to throughout all the lesson materials except in the case of the criterion tests. PT 5 is a criterion test and is similar to PT 3 except that no feedback or correction is given on the first presentation of the list. If the student meets criterion (i.e., no errors) on the first cycle, he leaves the word list presentation block. If he does not meet criterion, he returns to the task described in PT 2, branches past PT 3, and is presented with those items in PT 4 which he missed in PT 5. He then makes a second cycle through PT 5 with correction and optimization to one initial correct response for each item.

The problem route for these five problem types is represented by the flow diagram in Figure 3. Correct response routes are indicated by solid lines and remedial branches by dotted lines.

Insert Figure 3 about here

The optimization routine referred to in Figure 3 cycles the student through the list of response requests, repeating only those items which were responded to incorrectly until each item has received one initial correct response.

Remedial Materials for Word List Presentation Sections.

Early studies designed to test the feasibility of the lesson material indicated that the performance of a small percentage of the students was lower than anticipated on the word list materials. It was difficult at that time to specify all of the problems of these slow learning students since the learning difficulties were varied and complex. However, the learning problems can be grouped into four categories:

1) A need for finer discrimination of initial or the final units of stimulus words either as letters or as sound patterns or both.
2) A need for practice with fewer words. The memory processing capacity of some students seemed to be over-taxed.
3) A need for extensive practice on given words that troubled a particular child.
4) A need for a richer semantic context within which to relate the new words being learned.

Four sets of remedial materials to supplement the word presentation section were developed to give additional specific practice to those students experiencing one or more of the above problems. The four remedial sets are called: 1) Letter-Word Saliency, 2) Alliteration and Rhyming, 3) Rich Sentences, and 4) Reduplication and Sentence Initiator.

The four blocks of remedial word presentation materials are sequenced after PT 1 in the word list presentation block. It will be recalled that in a conventional sequence for the word list presentation block in PT 2 (Figure 3),

Insert Figure 4 about here

the picture is presented on the projector without the printed word, and a touching response to the appropriate word on the scope is requested. A counter is set at the beginning of PT 2 so that it tallies the errors for each initial presentation cycle but not for the optimization routines. At the end of the first cycle, the counter is examined. If its contents exceed a predetermined value (e.g., 50 percent of the number of response requests) the student is branched to the first assigned block of remedial material.

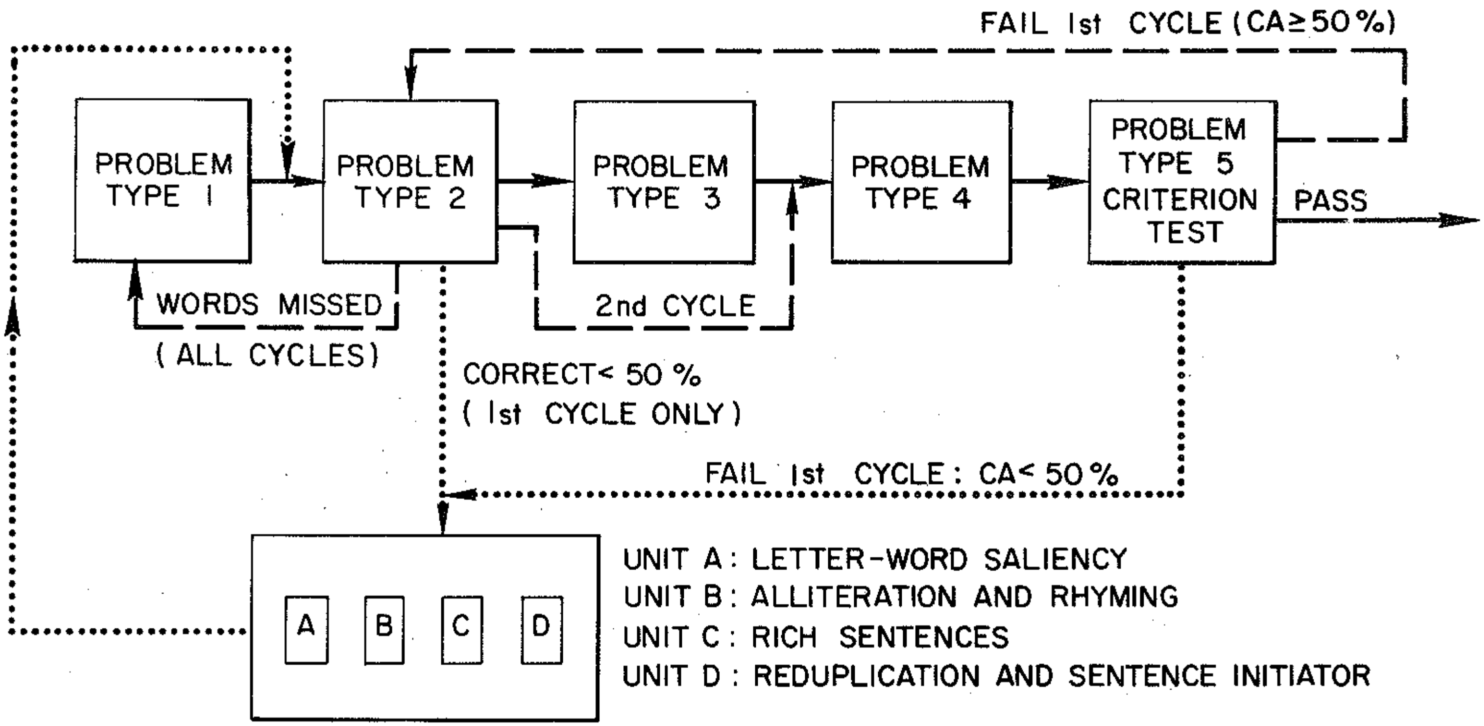

Figure 4.

Flow Diagram for Word List Block with Remedial Units

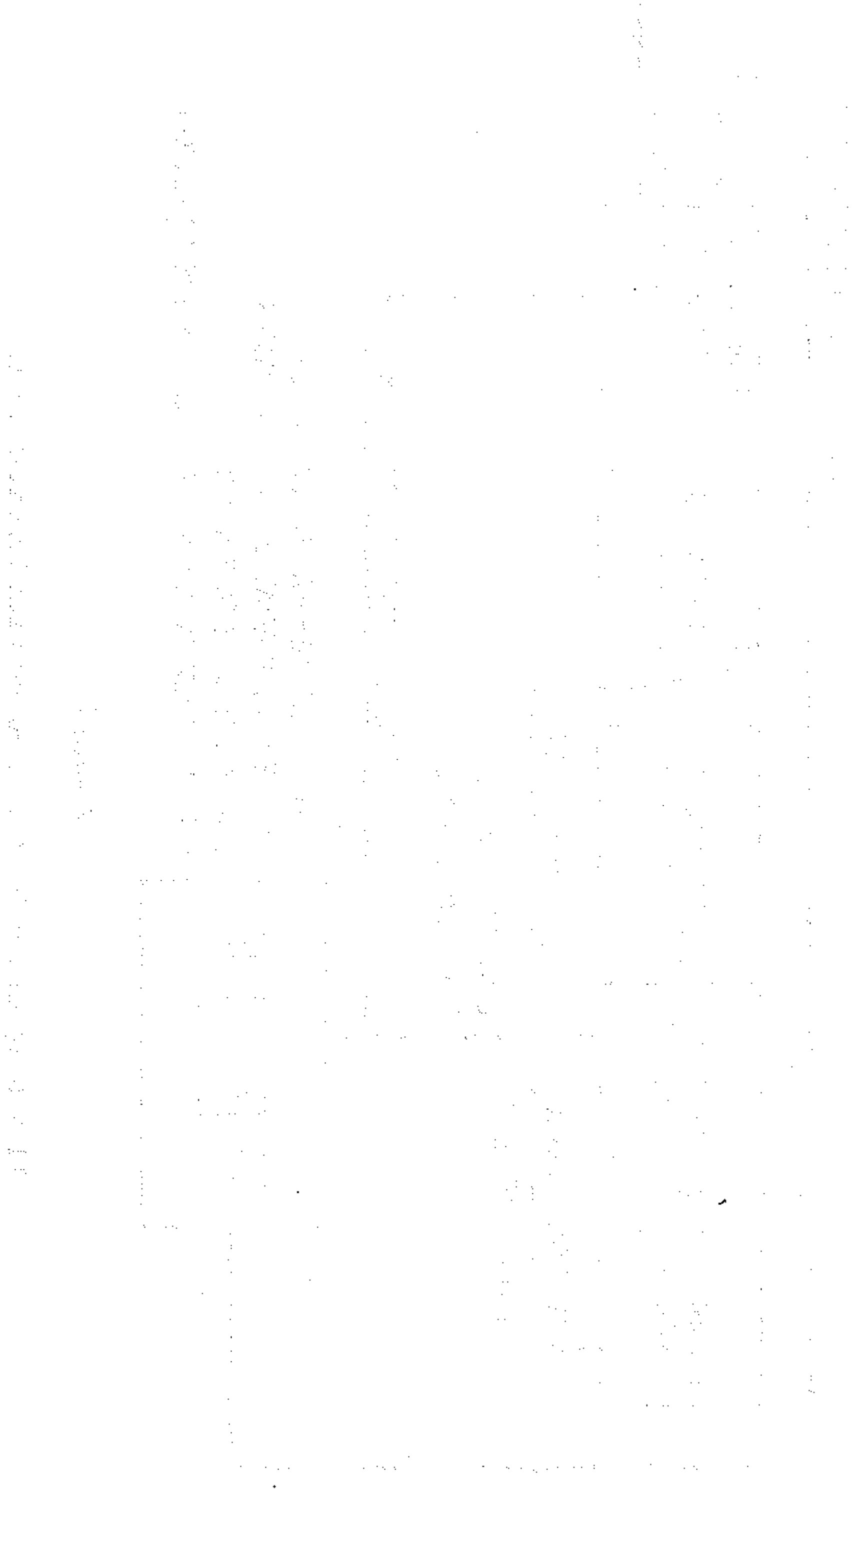

After completion of the remedial block, the student is returned for another cycle through PT 2. The error counter is again set to zero and the above operation is repeated. If the student again has an error rate of greater than 50 percent, not including the optimization routine, he is branched to the second assigned block of remedial material. The branching process is repeated if necessary until all four blocks of remedial material have been exhausted. In the case where a student completes a fifth failing cycle in PT 2, he is branched out of the instructional material and the proctor is automatically called. It is possible then for a student to receive from zero to four remedial presentations according to the evaluation after a cycle through PT 2.

If a student meets the criterion for Problem Type 2, he proceeds through Problem Types 3, 4 and 5, to the word criterion test. The decision counter is again set to zero at the beginning of the word criterion test. If a student meets criterion he proceeds to the next section of lesson material. If his performance in the test does not meet criterion and his error rate is less than 50 percent of the number of response requests, he is branched to PT 2. If, however, the error rate exceeds 50 percent, he is branched directly to the next of the remaining remedial sections before being recycled through PT 2, 4 and 5. In a second cycle through PT 5, the student is given an optimization routine and allowed to proceed to the next section of the curriculum. Therefore, there are two ways by which a student can reach the remedial material. The first is to exceed the error criterion for PT 2 and the second is to exceed the error criterion during the first cycle of PT 5.

The same general procedures apply to either method of entering the remedial material. It must be noted, however, that only one entrance to the remedial material may be made from the word list criterion test since the second pass through the criterion test is purely an optimization routine with no error counter being set for second branching to remedial loops. An evaluation of the remedial material and the remedial sequence will be made in the following manner. For each word list presentation section (that is to say, for each lesson), the four blocks of remedial material have been randomly ordered. The response data of students receiving remedial material will be subjected to off-line analysis in an attempt to discover if the increments in performance following each of the remedial blocks differ significantly from each other.

Matrix Construction.

Certain rhyming and alliterative word patterns can be considered regular; that is, certain graphic sequences correspond without exception to certain sound sequences. Fries (1963) and Bloomfield (1942), among others, have suggested that such patterned word regularity may be a key concept in learning to decode English print. This decoding concept is implemented in the Stanford materials via the sequencing of monosyllabic patterns as shown in the Vocabulary Sequence Chart (Table 1). The sounding matrix is an instructional technique which allows for practice in learning to associate orthographically similar sequences with appropriate rhyme and alliteration patterns.

Insert Table 1 about here

Table 1.

LESSON SEQUENCE

LEVELS	vc	cvc	ccvc	cVc~~e~~ ccVc~~e~~ cccVc~~e~~	cV ccV	cvcc ccvcc	cvvc ccvvc	cvv ccvv cccvv	cccvc cvccc
I 13 lessons	ac	cac							
II 19 lessons	ic	cic	ccac						
III 23 lessons	ec	cec	ccic	cAc~~e~~ ccAc~~e~~	cA				
IV 29 lessons	See Note 2.								
V 23 lessons	oc	coc	ccec	cIc~~e~~ ccIc~~e~~	cI cY cE ccY ccE	cacc			
VI 36 lessons	See Note 2.						ai oo ight	ay	
VII 43 lessons	uc	cuc	ccoc	cEc~~e~~ cOc~~e~~ ccOc~~e~~ cUc~~e~~ ccUc~~e~~	cO ccO	cicc cecc ccacc ccicc ccecc			a ccc i c a e c i ccc e
VIII 65 lessons	See Note 2.		ccuc			cucc cocc ccucc ccocc	ow oy ee ie ea ei ew		o cccuc cuccc o

Note 1: c = any consonant; v = any short vowel; V = any long vowel.
Note 2: Less frequent and less regular variations of preceding patterns (e.g., post-vocalic r, w, etc.)

Rhyming patterns are presented in the columns of the sounding matrix.

	an
r	ran
f	fan
c	can

Alliteration patterns are presented in the rows of the matrix.

	an	at	ag
r	ran	rat	rag

The matrix is constructed one cell at a time. The initial consonant of a consonant-vowel-consonant (CVC) word is termed the initial unit and the vowel and the final consonant are termed the final unit. This division is maintained later when the words are composed of consonant clusters and diphthongs. The intersection of an initial unit row and a final unit column determines the entry in any cell.

The problem format for the construction of each cell is divided into four parts: Parts A and D are standard instructional sections and Parts B and C are remedial sections. The flow diagram (Figure 5) indicates that remedial Parts B and C are branches from Part A and are presented independently or in combination.

Insert Figure 5 about here

The following example of a matrix cell construction is presented in some detail in order to highlight the instructional sequence and remedial branches.

The student sees the empty cell with its associated initial and final units and an array of response choices. He hears the audio message indicated by response request number 1 (RR 1) in Figure 6.

Insert Figure 6 about here

If the student makes the correct response (CA) (i.e., touches ran) he proceeds to Part D where he sees the word written in the cell and receives one additional practice trial (Figure 7).

Insert Figure 7 about here

In the course of an errorless trial the student hears the word pronounced three times and he must identify and pronounce it twice.

In the initial presentation in Part A, the multiple choice items are designed to identify three possible classes of errors:

1) The initial unit is correctly identified but the final unit is not.
2) The final unit is correctly identified but the initial unit is not.
3) Neither the initial unit nor the final unit is correctly identified.

If, in Part A, the student responds with fan he is branched to remedial section B where attention is focused on the initial unit of the cell (Figure 8).

Insert Figure 8 about here

If a correct response is made, the student is returned to Part A of the problem for a second attempt which will be optimized for a correct response. If an error response (WA) is made in Part B, the indicator is displayed above the initial unit beside the cell at the point indicated by the first circled superscript 1 in the audio message and held until the second superscript 1'.

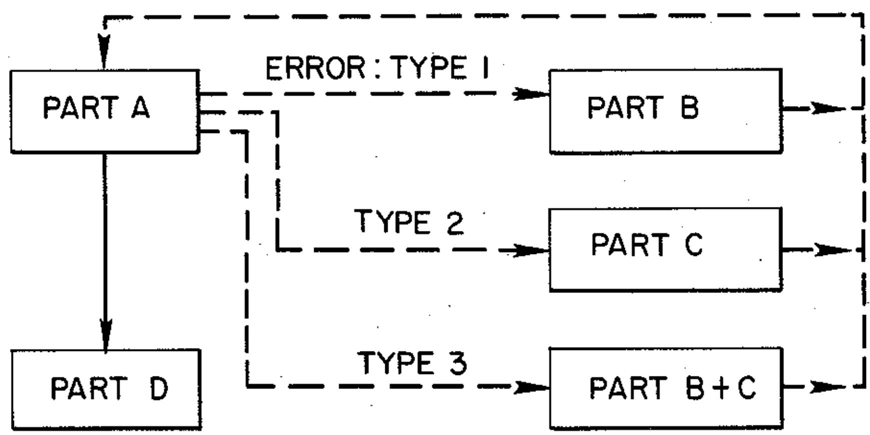

Figure 5.

Matrix Flow Diagram

	an
r	$\boxed{}$

rat
bat
fan
ran

RR 1: Touch and say the word that belongs in the empty cell.

CA: (Branch to Part D)

WA 1: No, rat = final → C, → A
fan = initial → B, → A
bat = other → B, → C, → A

↓CA
WA 2: No, touch and say ran.

Figure 6.

Matrix Problem, Part A

	an
r	$\boxed{\text{ran}}$

RR 1: Good, you have put ran in the cell. Touch and say ran.

CA: Good, ran. (→ next problem)

↓CA
WA: No, touch and say ran.

Figure 7.

Matrix Problem, Part D

At the completion of the response request, the indicator is placed beside the correct response area and held until a correct response is made (CA).

If, in Part A, the student responds with rat, he is branched to remedial section C where additional instruction is given on the final unit of the cell (Figure 9).

Insert Figure 9 about here

The procedures in Part C are similar to those of Part B. It should be noted in the remedial section that the initial letter is never pronounced (Part B) whereas the final unit is always pronounced (Part C).

If, in Part A, the student responds with bat, he is branched through both initial and final unit remedial work.

When the student returns to Part A after completing a remedial section, a correct response will advance him to Part D as indicated. If a wrong answer response is made on the second pass, the indicator arrow is placed beside the correct response area and held until a correct response is made. If the next response is still an error, a message is sent to the proctor and the sequence is repeated from the beginning.

When a student has made a correct response on Parts A and D, he is advanced to the next word cell of a matrix which has a problem format and sequence identical to that just described.

The individual cell building is continued until the matrix is complete. The majority of the matrices in the lesson material contain from six to twelve words and nonsense syllables. Nonsense words are considered legitimate cell entries under the following constraints:

1) Nonsense items are occurrent English syllables.

2) Nonsense items are not used which represent "regular" but unconventional spellings for common monosyllabic words. For example, although sed represents a regular spelling for the initial English syllable in words such as sediment, it would not be presented in matrix format since it would be considered an unacceptable spelling for the homophonous monosyllabic word said.

3) The proportion of nonsense words in the matrix is less than 40 percent of the total cell entries.

When the matrix is complete, the entries are reordered and a criterion test over all cell entries is given (Figure 10).

Insert Figure 10 about here

Randomized requests are made to the student to identify the cell entries. Since the first pass through the criterion matrix is a test trial, no reinforcement or correction is given. Errors are categorized as initial, final, and other. Remedial exercises are provided for both initial and final errors. The branching procedure is similar to that in the cell construction section with the exception that, whereas the branching in the cell construction block is contingent on each separate response, the branching in the criterion test is contingent on the total performance in the test. If the percentage of total errors exceeds some specified criterion percentage (e.g., 20 percent) of the total responses, the category registers are examined. If all the errors are recorded in one category (initial or final), only the remedial material appropriate to that category is presented. If the errors are distributed over both categories both types of remedial material are presented. This material is highly similar to conventional phonics exercises.

an

r $\boxed{}$

f
r
d

RR 1: Touch the initial unit of the empty cell.

CA: Good.

WA: 1 (r $\boxed{}$) No, this is the initial unit of the cell 1' so touch this.

CA $\rightarrow$ $\begin{pmatrix} f \\ r \\ d \end{pmatrix}$

Figure 8.

Matrix Problem, Part B

an

r $\boxed{r}$

an
at
ag

RR 1: Touch and say the final unit of the cell.

CA: Good.

WA: 1 ($\overset{an}{\boxed{}}$) No, an is the final unit of the cell 1' so touch and say

CA $\rightarrow$ $\begin{pmatrix} an \\ at \\ ag \end{pmatrix}$

Figure 9.

Matrix Problem, Part C

	at	an	ag
f	fat	fan	fag
r	rat	ran	rag
c	cat	can	cag

Touch and say:

RR 1: ran

RR 2: cag

RR 3: rat

⋮

Figure 10.

Matrix Criterion Test

The standard optimization scheme in which an initial correct response is obtained for each item is used for all sections of remedial material.

After working through one or both of the remedial sections, the student is looped back for a second pass through the criterion matrix. The second pass is a teaching trial as opposed to the initial test cycle; the student proceeds with the standard correction and optimization routines.

Compound Words.

The approach to compound words assumes the existence of a learning-transfer process in which the child knows how to read one of the two elements that form the compound word. The learning task consists of reviewing this known word and learning the unknown part of the compound. In addition the child must comprehend the conventional meaning of the compound word as well as its role in a fairly rich sentence. Thus we are utilizing compound words to study the degree to which children can transfer partial reading mastery (i.e., recognition of one of the two words in the compound) to a variety of reading contexts.

Compound words are introduced in Level III and are initially composed of two known monosyllables (e.g., bat and man are mastered prior to the presentation of the compound batman). Sequences are then introduced comprised of five compound words of which only one of the elements is known. Words are selected according to the following three criteria:

1) frequency in initial reading materials,
2) imaginative possibilities so that semantically rich context sentences may be written,
3) the opportunity to vary the known word in initial and final position in the five compound words (e.g., hatbox, firehat, hatband, etc.).

Polysyllabic Words.

Our approach to non-compound polysyllabic words focuses primarily on the role of stressed syllables in their relationship to verbal rehearsal and memory processing. We hypothesize that polysyllabic words with stress on the first syllable and a reduced or neutral vowel (usually transcribed as schwa (ə)) in the final syllable will be most easily learned. In the Stanford Reading Curriculum the acquisition of polysyllabic words is, as with compound words, a transfer learning process in which the first syllable of a polysyllabic word follows the pattern previously taught in the word list and decoding exercises. The stress and the number of syllables involved in the sequence are tightly controlled. The precise rules for sequencing polysyllabics are as follows:

1) the words are, at first, bisyllabic,
2) the first syllable of the word is stressed,
3) the vowel of the final syllable is represented as phonemic /ə/, excepting /iŋ/, /liy/, /iy/,
4) the first syllable is regular in that it follows the pattern learned in the word list and matrix exercises,
5) the word is monomorphemic where possible, and
6) the words are found in the Lorge-Thorndike Word List (1944).

Utilizing these rules, the unstressed second syllables shown in Table 2 were added to an appropriate vocabulary item to form a polysyllabic word.

Insert Table 2 about here

Polysyllabic words are introduced in Level IV. First presented are double rhymes in which the list words for a given lesson are both graphemically and

Table 2.

Unstressed Second Syllables

Initial Syllable + $\left\{ \begin{matrix} ə\mathrm{C} \\ \mathrm{C}ə\mathrm{C} \end{matrix} \right\}$				Stem + V	Stem + CV	Stem + CəC
age	el	ic	on	y	ly	let
al	le	ing	or			ling
ar	en	ip	ot			ness
	er	it	ow			
	et					

phonemically alike except for the initial consonant (e.g., bitter, litter, etc.). Subsequently, words are presented in which the second syllables rhyme and represent morphemically similar elements (sapling, starling). In general, polysyllabic words that are of high frequency and are useful for the story material were selected. We assume the words are understandable to the children, although no semantic criteria were specified.

The sequence of polysyllabic vocabulary allows for a build-up from two to three syllable words. There are direct systematic variations given between monomorphemic and polymorphemic word entries. All of these variations will be evaluated for their influence on the learning process.

The lesson structure for polysyllabic words consists of a sequence of approximately five words. For each pattern type (double rhyme, second syllable rhyme, morphemic or non-morphemic second syllable) the child is given a screening test similar to the matrix screening test. The child is asked to touch the cell where the two elements joined together such as <u>lad</u> + <u>der</u> forming <u>ladder</u> belong. Should the child be unsuccessful, he is presented with a rhyme utilizing the words from the screening test, followed by simple phonemic-graphemic decoding exercises adjusted to that particular pattern of polysyllabics. The child performing successfully on the screening test, and the child who has been presented with the phonemic-graphemic exercises then proceed to a lesson utilizing two words from the screening test and three to four others which follow the pattern being presented.

<u>Sentence Initiators.</u>

In reading written discourse, our learning objective is to train the initial readers to pronounce sentences with conventional intonation. A common reading behavior for beginners is giving "final word stress" to each of the words in a sentence. We wish to teach the child to adjust his timing, pitch, and stress contours so as to read sentences with intonation patterns commonly found in speech. This learning task takes on significance when one becomes aware that our CAI system does not include a speech analysis device nor is such a device even technologically available at this time. We depend upon our initial reader to be his own speech analyzer in this section of the Stanford Reading Lessons.

We do have available one of the important features of any CAI system, namely, the precise control of the timing of all learning events. The virtue and logistic nuisance of a CAI system is the required specification of all the time intervals for each of the instructional and reinforcement events within every problem. In this, as well as other sections of the lessons, we utilize the pacing feature of the CAI system to accelerate the reading-speech output of the children and, consequently, cause the children to approximate the speed and rhythms of normal speech while reading.

We maximized the expectancy of conventional sentence intonation by selecting high frequency sentence components such as "It's a," "That's a," "They can," etc., that are utilized commonly in the introduction of high frequency nouns and verbs in speech. These high frequency initiating sentential strings were selected from Carterette's (1965) list of multi-word units uttered by six-year old children during a free discussion session. It is to be noted that these word strings are pronounced by the children as if they were polysyllabic words. To each sentence initiator we attach a noun or a verb the children have previously learned in the list and matrix sections of the lessons. This provides us with an opportunity to review these words.

After a number of pilot investigations focused on the timing of the learning events, we found the following experimental procedure to be most promising in terms of achieving our objective. The sequence of the procedure is as follows:

1) anticipation interval; a sentence appears on a display device and the child is given two seconds with which to attempt an oral reading.
2) reinforcement interval; the audio device then plays a reading of the sentence to the child.
3) modeling interval; the child is given two seconds to repeat the reading of the sentence. For each new sentence the cycle is repeated.

In reinforcement terms, the child is expected to use self-evaluation in order to note any discrepancies between his reading and the machine's reading.

Each exercise in the "sentence initiator" section contains 25 sentences. Each initiator is presented approximately five to ten times in a given lesson and for a total of 100 to 125 exposures in all the lessons. At this point, the sentence initiator is utilized in the story material.

The manner of presentation of the sentence initiators allows the student to monitor his own production by comparing it with the audio model. Additional monitoring of the student's performance in the sentence initiator material is accomplished by the proctors and by sample recordings made on a separate recording system.

2. Comprehension

Producing the appropriate verbalization when confronted with some orthography (decoding or word-attack skills) is a necessary component but not a complete definition of reading. An equal and ultimately more important aspect of reading may best be termed "comprehension." Although a great deal has been written about "reading comprehension," it is not at all clear what we mean by the term. The general approach in current reading materials and reading achievement tests is at the paragraph level and primarily employs practice in recall or identification of specific details and sequence of events and in identifying the "main idea."

While using the standard techniques, we have also attempted to look at the question at the sentence level. Although we make no claim of achieving a complete definition of the field, we advance three propositions as necessary components of such a definition. In order to maintain that one comprehends a sentence one must be able to demonstrate that one has: 1) an appropriate set of semantic associations for each lexical item in the sentence, 2) an <u>operational</u> knowledge of syntax, and 3) the ability to identify the lexical items in the sentence which convey a given piece of information.

The next three sections of lesson material to be described are designed primarily to gather data pertinent to the above propositions. They are not, in the strict sense, teaching sections except that, as is the case throughout the lesson material, an overt correct response is required for each problem before the next one is presented. We do not feel that enough is known about comprehension to justify elaborate branching, remediation and optimization routines. Hopefully, some indications of how we might develop reasonable teaching strategies in this area may be found in the data collected in the following sections of the lesson material.

<u>Usage</u>

A list of words is displayed on the scope. Definitions are given auditorally one at a time and the student must identify the word that matches each definition by a touch-probe response. The definitions are chosen with the following constraints:

1) If the word appears in the Rainbow Dictionary (Wright, 1959), all the meanings defined in that dictionary must be used.

2) If the word does not appear in the Rainbow Dictionary but does appear in the Thorndike-Barnhardt Beginners' Dictionary (1964), at least one of the definitions must be used. The choice of usages is governed by the criteria of frequency and usefulness in story material.

3) If the word does not appear in either dictionary, it will not be included in the usage section nor used in succeeding lesson materials (e.g., sentences, stories).

A strict "dictionary definition" format is avoided in defining word items. Standard definitions have been re-constructed to stress functional meanings. This re-construction can be seen in the following example. For the word <u>bat</u> one dictionary definition is: "a stout wooden stick or club, used to hit the ball in baseball, cricket, etc." In the lesson materials this is reformulated in the following instruction: "Touch and say the word that means something you might use to hit a baseball."

A correct response is reinforced by confirmation of the functional definition, "Yes, you hit a baseball with a bat." An incorrect response is corrected ↓CA in a straightforward manner, "No, you hit a baseball with a bat. Touch and say bat." Thus, for an initial correct response the definition is associated with the word twice. For an initial incorrect response, the definition is associated with the word three times since the student will receive the reinforcing confirmation after his correct response has been made. No branching or optimization routines are employed in the usage section.

<u>Form Class</u>.

As part of the concept of "comprehension" of a sentence, one must consider the child's basic understanding of English syntax. One behavioral manifestation of a child's syntactic sophistication is his ability to functionally group words into common form classes (nouns, verbs, modifiers, etc.). This section provides lesson materials that both assess and teach the form class characteristics for the words just presented in the matrix section.

The following type of problem is presented to the student:

Dan saw the	tan fat man run	hat.

Only one of the words in the column will make sense in the sentence. Touch and say the word that belongs in the sentence.

CA: Yes, Dan saw the tan hat.
Do the next one.
↓CA

WA: No, tan is the word that makes sense. Dan saw the tan hat.
Touch and say tan.

The sentence is composed of words that are in the well-practiced reading vocabulary of the student (i.e., they have been presented in previous or current lessons). The elements of the multiple choice set include a word which is of the correct form class but is semantically inappropriate, two words that are of the wrong form class, and of course the correct word. A controlled variety of sentence types is employed and the answer sets are distributed over all syntactic slots within each sentence type.

Response choices are categorized in rather broad terms as <u>nouns</u>, <u>verbs</u>, <u>modifiers</u>, and <u>other</u>. The response data are examined for systematic errors over a large number of items. Examples of the kinds of questions that will be asked of this data are:

1) Are errors for various form classes in various sentence positions similarly distributed for a given student? for the general population?

2) How are response latencies affected by the syntactic and serial position of the response set within the sentence? by the sentence? by the sentence length?

Answers to these questions and others of a similarly general nature should provide information that will permit more systematic study of the relationship of sentence structure to reading instruction.

Inquiries.

Individual words in sentences constitute unique and conversationally correct answers to questions. These questions take the interrogative form "Who - - -?, What - - -?, How - - -?" etc. The ability to select the word in a sentence that uniquely answers one of these questions demonstrates one form of reading comprehension at the level of the sentence. The set of exercises described in this section constitutes an assessment of this reading comprehension ability.

John hit the ball.

Touch and say the word that answers the question.

RR 1 Who hit the ball?

CA: Yes, the word "John" tells us who hit the ball.
↓CA
WA: No, John tells us who hit the ball. Touch and say John.

RR 2 What did John hit?

CA: Yes, the word "ball" tell us what John hit.
↓CA
WA: No, ball tells us what John hit. Touch and say ball.

As in the form class section, each sentence is composed of words from the student's reading vocabulary. A wide variety of sentence structures are utilized beginning with simple subject-verb-object sentences and progressing to structures of increasing complexity. The data will be examined for effects caused by the different sentence types.

Stories Read to the Children.

From the very beginning lesson the children have an opportunity to listen and react to children's literature. This section of the lesson material has been designed for the following purposes:

1) to expand and widen the children's literary experience,
2) to indicate how the rhythm and flow of language is consistent with the thematic material,
3) to improve the children's listening comprehension in order to compar it in future points with the children's reading comprehension,
4) to receive the children's reactions to such variables as
    a) thematic content of the story,
    b) length,
    c) style (narrative versus expository form).

In the actual lesson format, the story is presented on the CRT, sentence by sentence, so the child can follow the print as it is read to him. Pictures are presented with stories if the Project staff believes the children need help in understanding certain concepts being presented. The mode of presentation simulates, in some sense, a parent holding the child in his lap and reading a story aloud -- an experience which unfortunately many children have missed.

At the completion of each story, the following types of questions are asked:

1) subjective ratings as to positive or negative reactions,
2) specific questions on factual identification of characters or events in the story,
3) identification of the main theme of the story,
4) identification of the sequencing of events.

In a given story only some of these questions are employed; we, in essence, let the content of the story determine the type of questions asked. The number of questions also varies although typically there are approximately five questions per story.

Stories the Children Read.

When the children have mastered a sufficient amount of vocabulary, including function and syntax words, stories which are actually read by the children are introduced into the reading lessons. These original children's stories are developed by staff members who are skilled writers and understand children's literature. The principal problem encountered in this process is that of creating stories around a highly restricted vocabulary. Therefore the following criteria were utilized in preparing and editing these original stories:

1) the variety of literary themes,
2) the children's familiarity with the thematic content and characters,
3) logical sequence of events,
4) simplicity of plot and characterization
5) appropriateness in the use of the restricted vocabulary,
6) story length.

These stories are first inserted in the lessons at Level IV.

The stories are presented on the scope in the same manner as a story is presented in a story book. The child is to read the story by himself. If he does not know a word, he is directed to touch the word and the entire sentence in which the word appears is read to him. Illustrations are used with the stories as they are needed at certain points to help develop an understanding of the concepts being presented. These illustrations were specified by the authors and prepared by the Project's art staff.

After the child has completed the story, he is presented with a series of questions to determine his level of comprehension. The types of questions are divided into four main categories:

1) questions that deal with the direct recall of facts such as the identification of characters and the sequence of events,
2) generalizations about main ideas which relate both characters and events within the story in which these characters and/or events are not presented in close contiguity,
3) inferential questions which require the child to relate information presented in the story to information stored in his memory about his own experience,
4) subjective questions which include personal ratings and opinions of the stories.

The kinds of questions asked about each story are determined by its content, therefore not all stories have exactly the same type of questions. Question types are carefully sequenced throughout the stories to provide equal representation. The number of questions asked about each story varies from 3 to 6 with generally 4 questions used. Again the questions and the related stories represent an attempt by the Project staff to collect appropriate information which will help us in formulating hypotheses about reading comprehension. Any theoretical development in this area must of necessity relate both the child's immediate information gathering ability and his ability to logically relate this new information with concepts stored in his memory system.

3. Games and Other Motivational Devices

Rhymes.

Rhymes are sequenced into a lesson as a listening activity to help the child develop competency in the discrimination of the rhyming and alliterative sounds of words and to demonstrate to the child the rhythmic use of language. The selection of rhymes for each lesson is based upon the sound patterns found in the matrix section of the lesson.

Games.

Games are sequenced into each lesson primarily to encourage continued attention to the lesson materials. The games are similar to those played in the classroom and utilize the terminology common to game-like situations such as baseball or bingo. In addition, each game has been structured to reveal any developing linguistic competency on the part of the child. For example, one game is centered on the child's ability to identify a given string of letters when presented in a concept-identification paradigm. The question of how children come to see orthographic commonalities is relatively unexplored, and the game is intended to provide us with some basic information with which we can more precisely direct our thinking. Moreover, the games give the children added practice on the sequenced vocabulary.

As indicated earlier, all of these games are intended to provide the children with intrinsic motivation to continue practicing much of the same material that has previously been presented. Ordinarily the words found in the games were introduced in the preceding lesson. Again it is to be noted that the responses in the game sections can be analyzed just as paired-associate items or as concept-identification items in more controlled psychological experimentation.

4. Review Lessons

Two types of review are to be found in the reading lessons. A continuous review is inherent in the learning materials which are sequentially introduced in a given lesson and subsequently incorporated in the activities of succeeding lessons. For example, words which are introduced in list and matrix exercises in Lesson N are reviewed in the sentence initiator and story materials of Lesson N+1. A more conventional type of review is also furnished by lessons designated as review lessons. These appear approximately every seventh lesson. Vocabulary and concepts previously introduced are re-presented but in different formats.

The nature of the review lessons varies as the vocabulary and skills of the student progressively develop. In Level I the review lessons consist primarily of straightforward matrix review; that is, the words which have been introduced in the preceding five to seven matrices are reordered and presented in new matrix formats. The review lessons of Level II focus primarily on the recently acquired vocabulary items, prepositions, and the inflectional concepts of the plural-s, third person singular-s and the -ing suffix.

In Level III the range of activities in the review lessons is considerably increased and basically furnishes the format for review lessons in the succeeding levels. Exercises on phonetic discrimination, word meaning, form class, English word order, rhyming exercises, analysis and synthesis of words, picture-sentence comprehension, compound and polysyllabic words, and verb forms are included.

PRODUCTION

Lesson Writing.

Overall Design. An overall design (scope and sequence) for the decoding aspect of the curriculum was developed during the first year of the project based on the psycholinguistic propositions stated above. The primary responsibility for the overall design was assumed by the principle investigator, the senior research associate and the staff linguist. It must be emphasized, however, that at every stage in the development of the curriculum the intuitions, experience and expertise of the entire staff have played a major role.

A well defined sequence of presentation of linguistic patterns is a necessary preliminary to actual lesson writing. The development of such a sequence, however, is a dynamic process which involves a complex interaction of study, deliberation, experimentation, and evaluation. The sequence as seen in Table 1 was not completed before production of lesson materials for the early levels began, nor is it to be considered complete in its present form. The specification of patterns beyond Level VIII is a continuing aspect of the current curriculum development effort.

Within-Level Sequencing. Since the sounding matrix determines the monosyllabic patterns, the vocabulary available for each lesson and the number of lessons to be contained in each level, the next step in lesson writing, after the overall-sequence has been established, is to generate and sequence the sounding matrices within a level. The responsibility for this task is usually assumed by a single writer although in the larger levels two writers may work as a team.

All possible words and nonsense words are generated for the level by permuting the consonants and consonant clusters about vowels under the constraint of the patterns and letters assigned to a given level. The resulting set of words and nonsense words are then sorted into matrices by a trial and error method. The goal is to produce an optimal set of matrices which give adequate practice for all the patterns. Each matrix must conform to the following specifications:

1) No matrix may exceed a size of 16 cells.
2) No matrix may contain a ratio of nonsense to real words in excess of 40%.
3) Socially unacceptable words must be excluded.
4) Nonsense words which are misspellings of common monosyllabic words must be excluded unless they form highly productive syllables in polysyllabic words.
5) Non-English nonsense words are excluded.

An exhaustive treatment of the patterns has been made in Levels I and II. In the later levels, however, the goal has been to provide adequate exposure and practice for each pattern and to withhold an increasing proportion of the possible words for a given pattern from direct teaching in order to use them as transfer items.

After the matrix generating writer or team has produced a reasonable set of matrices for a level, the result is reviewed by the entire staff and suggestions are made for improving the sequence. The matrix sequencing process may go through several iterations of trial and review before a consensus is reached and a set of matrices is approved for a level.

<u>Lesson Level</u>. Once the sounding matrices have been generated and sequenced for a given level the actual lesson writing begins. Each lesson which will contain a matrix is called a "standard lesson," and has many learning tasks and problem types. A discussion of the various standard problem blocks will be found in the section headed "Curriculum Description." Each writer is responsible for the production of some set of problem blocks within a level. A balance is sought between the more mechanical blocks (e.g., matrix construction) and the more creative blocks (e.g., usage, form class, poems) for each writer's assignment.

Writing a lesson is quite different than writing a portion of an ordinary text book. It is more akin, perhaps, to writing a movie or TV script. Everything the student is to hear, see, and do must be clearly defined for every minute he is at the terminal. Each display on the CRT must be unambigously specified and desired changes in the display must be cued precisely with the audio messages. Displays on the projector must be sketched in sufficient detail for communication with the artists. The timing of projector displays must also be clearly specified.

Although the branching logic, test criteria, optimization routines and within-lesson task sequence have been determined in general by staff consensus, the lesson writer must communicate all such information clearly to the curriculum programmers for each problem in every lesson. Lesson production forms have been developed for all the standard problem types to minimize the amount of repetitive detail that is required of the writers and to insure uniform presentation of many of the learning tasks for experimental purposes.

Non-standard lessons are those which present non-recurring or very infrequent learning tasks. The lesson which covers the comparative morphemes -er and -est is an example of a non-standard lesson. In all such material, whether it is an entire lesson or a section of an otherwise standard lesson, the writer has to explicate all of the information for each problem without the benefit of forms -- a tedious and time consuming task.

The writing staff is composed mainly of experienced primary teachers and graduate students in education. It is difficult, however, to define clearly what is meant by the term "writing staff" since most of the writers have other responsibilities within the project and nearly all members of the project staff, with the exception of the programmers and artists, have made major contributions to the writing effort at one time or another. It is even more difficult to make a reasonable estimate of the man-hours required to write a lesson. The production rate for lesson material, including editing but excluding audio recording, art production and programming is on the order of 100 lessons a year which represents something in excess of 90 man-hours per lesson.

Editing. Editing the lesson materials is a vital and sensitive process in the production of the Stanford CAI Reading Curriculum. Vocabulary constraints imposed by the careful sequencing of patterns must not be violated and the branching logic of each problem block must be checked for consistency. The lesson editor is also responsible for the overall quality of the lesson material in the sense of the interest and appropriateness of the sentences, poems and stories.

Several methods of organizing the editing process were tried during the early phases of the project, such as section sub-editing, and round-robin multiple editing. It was largely through these efforts to control and make manageable the extremely complex editorial requirements that the standardized formats were developed. The problem of vocabulary control became too difficult in the later levels to be met by the normal editorial procedures. A dictionary or word list was compiled which could be stored in a computer and continuously updated. A dictionary look-up program was written which matches the displayed text of each lesson against the stored dictionary and prints out the words used in the lesson not found in the dictionary up to that point. The resulting print-out is examined by the lesson editor to identify those words which are allowable transfer items. Corrections are then made in the lesson text which eliminate the remaining words detected by the computer program (i.e., those words not in the dictionary at that point in the pattern sequence and not usable under the rules for transfer items). The above procedure assures a high degree of consistency and accuracy in vocabulary control.

Audio Production.

Message Numbering. The first step in producing audio tapes is the numbering of the individual audio messages. The coders number each message in the lesson books after the lessons have received their final editing. The numbering of audio messages is necessary for two reasons. First, the coders refer to the messages only by number when coding the lessons. Secondly, the messages do not coincide completely with the response requests. As an example of the latter, the wrong answer message, as it appears in a lesson consists of two parts: the word "no" and a correction message. In the lesson coding this is broken into two messages. The second message is played after an overtime response since it is not appropriate to say "no" to a student who has done nothing. Both messages are played after a response which was in error.

As the messages are numbered, an indicator [X] is placed at each point in the message where a display change on the projector or CRT must be made in synchronization with the audio. These display changes most commonly involve the addition or deletion of an arrow or an underline to an existing display, hence the X's are referred to as emphasis indicators.

The messages are numbered in blocks of fifty to facilitate future editing or revisions. The message code number consists of an alphabetic character followed by two digits. The numbering scheme can uniquely identify 1300 messages for each lesson. In practice, the average lesson contains approximately 400 messages.

Script Preparation. After the messages have been numbered by the coders, an audio script for each lesson is produced by the typists from the lesson books. The audio script is necessary for speed and accuracy in recording. All irrelevant material in the lesson book, such as response identifiers and branching instructions, is ignored in the production of the audio script. The messages are typed in sequential order. The only material included in the audio script, besides the actual messages, are the message numbers and the emphasis indicators.

The scripts are carefully proofread for accuracy. Minor typographical errors such as substituting man for fan and misplaced emphasis indicators are difficult to detect except by team-proofing against the original lesson material. Such errors, if missed in the proofreading, are increasingly more difficult to correct at each succeeding step in the audio production processes.

Recording. The master narration tape is recorded on an Ampex PR-10 two-track tape recorder. The messages are recorded on the "A" channel using an LTV 1800 cardioid microphone. A constant 400 cycle tone is recorded on the "B" channel in conjunction with each message from a Hewlett-Packard 200 AB oscillator controlled by a tone key.

The requirements of the system, which will be discussed in detail in the section "Audio Compiler," specify that the 400 cycle tone must be generated by depressing the key at least 0.25 seconds before the message is started and continued until at least 0.25 seconds after the message is finished. Further, the intermessage gap must be at least 3 seconds and less than 10 seconds in duration and the emphasis indicator must be signaled by a cessation of the tone of at least 0.5 seconds and no more than 1.5 seconds.

Practical considerations of audio consumption at the student terminals require that the above minimum specifications be adhered to as nearly as possible. The recording narrators must not only exhibit pleasant voice quality, clear enunciation and good intonation patterns but also develop considerable skill in the use of the tone key.

The master narration tapes are produced in a sound-proof recording studio at Ventura Hall on the Stanford campus by two female narrators, each recording for four hours at a session. Experience has shown that a narrator cannot maintain accuracy and efficiency over longer periods.

Each lesson is concluded with a story which is read to the student while the orthography is displayed on the CRT. This section of the lesson, known as the "listening stories" is recorded by a group of professional actors from the Stanford Repertory Theater. Professional actors are utilized for the recording of this material to insure an interesting, dramatic, and appropriate presentation of the stories. It is assumed that such a presentation will maximize the motivational aspects of this section of the lesson material and at the same time provide an effective introduction to literature.

Since the orthography is displayed sentence by sentence, each sentence is considered as a separate message and no emphasis indicators are necessary to synchronize the CRT display and the audio. The actors then are only required to make the appropriate tone key manipulations between sentences, a technique which is quickly learned, and can therefore concentrate their attention on producing an artistic reading.

The master narration tape for each lesson is edited and corrected by the narrator immediately after recording. The tapes are then traded and given a final edit with each of the two narrators editing her colleague's tape. The major aspects of the narration tape which are carefully attended to in the editing process are: correctness of the messages, clarity of enunciation, proper tone signal overlap, and accuracy of the emphasis indicators. The rather elaborate and painstaking editorial procedure is a vital part of audio production since correction of tapes is very difficult once they have undergone the audio assembly process.

Audio Assembly. After a master narration tape has received the final editing by the narrators it is sent to the Brentwood laboratory for the generation of a master machine tape. A machine master is a four-track tape. The first two tracks contain the audio messages taken from the master narration tape. The third track is reserved for on-line student recording and the fourth track contains the tape segment addresses.

A machine master is produced by a computer program which reads the tone signals on the master narration tape, evaluates the length of a set of such signals and records their associated verbal messages on the machine master using an optimal packing strategy. Figure 11 is a graphic example of what is meant by "optimal packing."

Insert Figure 11 about here

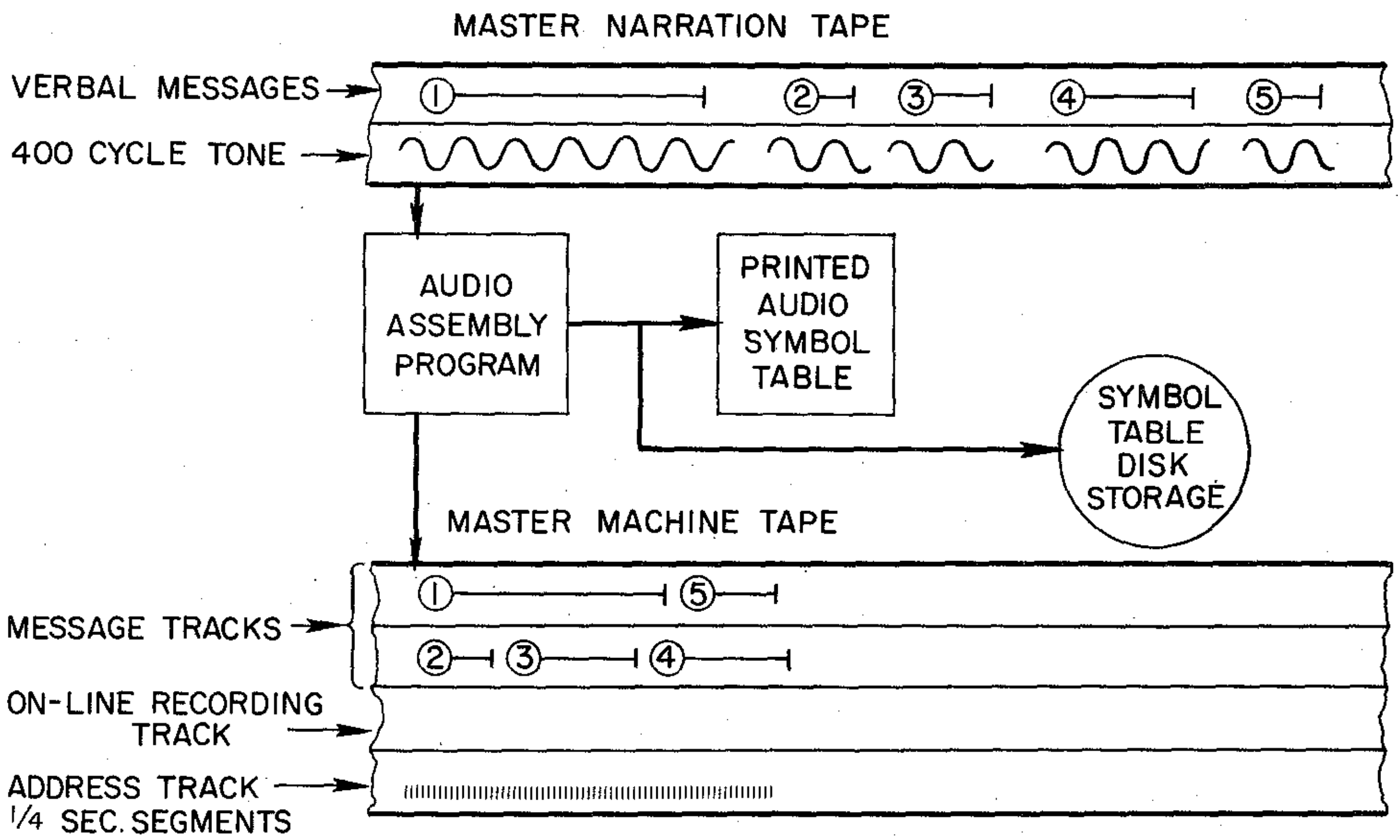

Figure 11.

Audio Assembly Routine

Verbal messages of varying length are recorded on the two message tracks in a manner which maximizes the number of messages per unit length of tape and minimizes the distance between any two messages. The latter condition also minimizes the search time required for positioning any message.

The result of the optimal packing strategy is a non-sequential ordering of messages on the machine master. To control an otherwise chaotic condition, a table is generated concurrently with the recording of the messages which specifies the location on the machine master (by track and segment number) of the beginning of each message, its sequential number on the master narration tape and its length. This table, called the <u>Audio Symbol Table</u>, is stored on the symbol table disk for each lesson. A printed audio symbol table is also produced for use by the programmers in the debugging process.

The complete process of producing a master machine tape and the associated tables is known as <u>audio assembly</u> and is carried out during production hours by a single computer operator, typically between midnight and 8 A.M.

<u>Tape Duplication</u>. The terminal tapes are produced from the master machine tape on a Viking 235 Duplicator. Three duplicates can be made simultaneously on the Viking at 7 1/2 inches per second. Tape duplication is under the control of a single technician and the duplication time follows an intermittent demand schedule since the process is independent of the 1500 system.

The terminal tape requirements for each day are determined by the distribution of students over the lesson material. After all students have completed a given lesson, all but three terminal tapes for that lesson are degaussed. The three remaining duplicates and the machine master are placed in storage. Three duplicates have been found to be sufficient to act as a buffer to absorb new students and for demonstration purposes.

Art Production. Production of the art work for the reading curriculum begins with the lesson writers who, in the course of writing the lessons, specify the necessary illustrations. The art specifications are made on the "projector" section of the visual pages of the lesson by means of rough sketches and descriptive notes. Illustrations are required for the word list sections, some sentence analysis sections, and all of the listening and reading stories and poems.

The amount of detail in the specifications necessary for communication between writer and artist varys widely for the different sections. For example, in the word-list section each illustration is well defined by the individual words and their sentence context. A minimum of specification is therefore required in the word-list section. On the other hand, the salient points of the story section which the author wishes to emphasize in each illustration must be clearly defined for the artists.

As in all phases of lesson production, the communication channels between the two groups, writers and artists, are open for a continuous two way flow of information. The illustration requests of the writers must occasionally be adjusted in terms of the artist's ability to visually convey abstract terms and the practical limitations of time and production facilities. Such compromises often require rewriting of individual frames. The early stages of art preparation for a block of lessons is therefore characterized by an open dialogue between the writers and the artists.

Art production requirements are currently being met by an art editor, one half-time and three full-time artists. The size of the art staff is determined by a compromise between long range requirements for illustrations and the space available to house the artists. In order to meet our total art commitments the present staff is required to produce an average of twenty plates a week per artist.

Techniques and media are dictated by the demands of high production and high quality. Most of the plates are done with felt tip marker on two-ply bristol board, with occasional use of colored ink, chalk, paste-up and tempera. An additional factor in the choice of media and palette is the complex set of requirements of color photo-reproduction.

A serious attempt has been made to include in the pictorial illustrations a balanced distribution of racial and ethnic figures as a reflection of the mixture of those elements in American society. The goal has been to feature all racial and ethnic types as central figures engaged in interesting activities, by means of artistically and socially honest and appropriate drawing.

Photo Reproduction.

The illustrations are photographed in blocks of lessons which are defined by the level structure of the overall sequence. The shooting is done at the Hospital Laboratory of the Stanford Photo Service with a 16 mm movie camera and single frame exposure.

Correct sequence of illustrations is vital, therefore the art editor, who originally assembled the plates, is present to assist the photographer at every shooting session. The plates are shot in blocks of lessons with 9 blank frames left between blocks to minimize the difficulty of splicing if later corrections or additions are necessary.

A work print is produced by the processing laboratory which is examined by the art editor. The work print is stepped through an IBM 1512 projector at the Brentwood Laboratory and is checked for proper sequence and acceptable color.

When a work print has received the final approval of the art editor it is returned to the processing laboratory and a release print is made. The release print, which contains pictures but no identification code, is combined, through a process known as A-B printing, with a code master to produce a master film. The master film is then used to produce the six terminal prints for debugging on the 1500 system.

The six terminal prints are mounted on self-threading, dust proof cartridges and are used, along with the audio tapes, in debugging coded lessons. When the debugging process is completed, a quantity of terminal prints are produced and mounted which will assure sufficient films and reserves to accommodate the maximum number of students anticipated to be working in a level at any given time.

Coding.

Before the lessons can be used in the CAI system they must be translated into a language which is understood by the computer. The first step in this process is accomplished by the lesson programmers by coding the lessons in Coursewriter II -- a CAI source language developed by IBM. A coded lesson is actually a series of Coursewriter II commands which cause the computer to display and manipulate text and graphics on the CRT, position and display film in the projector, position and play audio messages, accept and evaluate keyboard and light pen responses, update the performance record of each student and implement the branching logic of the lesson flow by means of manipulating and referencing a set of switches and counters. A typical lesson will require 9,000 Coursewriter II commands for its execution.

Table 3 shows the audio messages and film images required for two sample problems along with possible addresses on the audio tape and film strip. What

Table 3.

Audio Script and Film Images with Hypothetical Addresses

Audio information	
Address	Message
A01:	Touch and say the word that goes with the picture.
A02:	Good. Bag. Do the next one.
A03:	No.
A04:	The word that goes with the picture is bag. Touch and say bag.
A05:	Good. Card. Do the next one.
A06:	No.
A07:	The word that goes with the picture is card. Touch and say card.

Film strip	
Address	Picture
F01:	Picture of a bag.
F02:	Picture of a card.

follows are the computer commands required to present two problems described in Table 3, analyze the students' responses, and record them on the data record. The left column lists actual commands to the computer controlling the instruction. (Labels L1, L2, etc. in the column on the far left indicate points which can be branched to.) On the right is an explanation of the results of the execution of these commands The first problem is explained command by command; the second problem is explained only in outline.

Insert Table 3 about here

Commands	Explanation
Pr	Problem: Prepares machine for beginning of new problem.
LD 0/S1	Load: Loads 0 into the error switch (S1). The role of switches and counters will be explained later.
FP F01	Film Position: Displays frame F01 (picture of a bag).
DT 5,18/bat/	Display Text: Displays "bat" on line 5 starting in column 18 on the CRT.
DT 7,18/bag/	Displays "bag" on line 7 starting in column 18 on the CRT.
DT 9,18/rat/	Displays "rat" on line 9 starting in column 18 on the CRT.
AUP A01	Audio Play: Plays audio message #A01. "Touch and say the word that goes with the picture."
L1 EP 30/ABCD1	Enter and Process: Activates the light pen; specifies the time limit (30 seconds) and the problem identifier (ABDC1) that will be placed in the data record along with all responses to this problem. If a response is

Commands	Explanation
	made within the time limit the computer skips from this command down to the CA (correct answer comparison) command. If no response is made within the time limit, the commands immediately following the EP command are executed.
AD 1/C4	Add: Adds 1 to the overtime counter (C4).
LD 1/S1	Load: Loads 1 into the error switch (S1).
AUP A04	Audio Play: Plays message #A04. "The word that goes with the picture is bag. Touch and say bag."
DT 7,16/→/	Display Text: Displays "→" on line 7, column 16 (arrow pointing at "bag").
BR L1	Branch: Branches to command labeled L1. The computer will now do that command (EP) and continue from that point.
CA 1,7,3,18/C1	Correct Answer: Compares student's response with the area 1 line high starting on line 7, 3 columns wide starting in column 18 on the CRT. If his response falls within this area, it will be recorded in the data with the answer identifier C1. When a correct answer has been made, the commands from here down to WA (wrong answer comparison) are executed. Then the machine jumps ahead to the next PR. If the response does not fall in the correct area, the machine skips from this command down to the WA command.

Command	Explanation
BR L2/S1/1	Branch: Branches to command labeled L2 if the error switch (S1) is equal to 1.
AD 1/C1	Add: Adds 1 to the initial correct answer counter (C1).
L2 AUP A02	Audio Play: Plays audio message #A02. "Good. Bag. Do the next one."
WA 1,5,3,18/W1 WA 1,9,3,18/W2	Wrong Answer: These two commands compare the student response with the areas of the two wrong answers, that is, the area 1 line high starting on line 5, 3 colums wide starting in column 18, and the area 1 line high starting on line 9, 3 columns wide starting in column 18. If the response falls within one of these two areas, it will be recorded with the appropriate identifier (W1 or W2). When a defined wrong answer has been made, the commands from here down to UN (undefined answer) are executed. Then the computer goes back to the EP for this problem. If the response does not fall in one of the defined wrong answer areas, the machine skips from this command down to the UN command.
AD 1/C2	Add: Adds 1 to the defined wrong answer counter (C2).
L3 LD 1/S1	Load: Loads 1 into the error switch (S1).
AUP A03	Audio Play: Plays message #A03. "No."
AUP A04	Audio Play: Plays message #A04. "The word that goes with the picture is 'bag'. Touch and say bag."
DT 7, 16/→/	Display Text: Displays "→" on line 7, column 16.

Commands	Explanation
UN	Undefined Wrong Answer: If machine reaches this point in the program, the student has made neither a correct nor a defined wrong answer.
AD 1/C3	Add: Adds 1 to the undefined answer counter (C3).
BR L3	Branch: Branches to command labeled L3. (The same thing should be done for both UN and WA answers. This branch saves repeating the commands from L3 down to UN.)
PR	Prepares the machine for next (2nd) problem.
LD 0/S1 FP F02 DT 5,18/card/ DT 7,18/cart/ DT 9,18/hard/	These commands prepare the display for the 2nd problem. Notice the new film position and new words displayed. The student was told to "do the next one" when he finished the last problem so he needs no audio message to begin this.
L4 EP 30/ABCD2	Light pen is activated.
AD 1/C4 LD 1/S1 AUP A07 DT 5,16/→/ BR L4	These commands are done only if no response is made in the time limit of 30 seconds. Otherwise the machine skips to the CA command.
CA 1,5,4,18/C2	Compares response with correct answer area.

Commands	Explanation
BR L5/S1/1 AD 1/C1 L5 AUP A05	1 is added to the initial correct answer counter unless the error switch (S1) shows that an error has been made for this problem. The student is told he is correct and goes on to the next problem. These commands are executed only if a correct answer has been made.
WA 1,7,4,18/W3 WA 1,9,4,18/W4	Compare response with defined wrong answer.
AD 1/C2 L6 LD 1/S1 AUP A06 AUP A07 DT 5,16/→/	1 is added to the defined wrong answer area and the error switch (S1) is loaded with 1 to show that an error has been made on this problem. The student is told he is wrong and shown the correct answer and asked to touch it. These commands are executed only if a defined wrong answer has been made.
UN	An undefined response has been made if the machine reaches this command.
AD 1/C3 BR L6	1 is added to the undefined answer counter and we branch up to give the same audio, etc. as is given for the defined wrong answer.

Thirty counters that can be used to keep track of a student's performance are available to the lesson programmer. During the instructional flow the current values of these counters are used to make branching decisions regarding what stimulus materials are to be presented next. For example, if the correct-answer counter for a particular class of problems has a high value, the student may be branched ahead to more difficult topics, whereas for a low value he may be branched to remedial work. These counters can contain any number from 0 to 32,767. They are normally set at zero at the beginning of a course and added to when desired. For example, counter 4 (C4) was used to record overtimes; each time the time limit was exceeded one was added to counter 4 (AD 1/C4).

There are also 32 switches available to the instructor. A switch is either in the zero or one position. These are used to keep track of previous events. For example, at the beginning of a problem, zero is loaded into S1 (the "error" switch). This indicates that no error has yet been made on this problem. If the student makes an error on the problem, one is loaded into S1. Then, if a correct answer is made on his second try, the command can be branched around adding one to the initial correct answer counter because the error switch (S1) is equal to one.

There are many features of the CAI system that are not demonstrated by the simplified example presented here. The pattern of the problems may vary widely from this sample. At various points in a lesson, criteria may be set which, if not met, may cause the student to branch to remedial problems or have the proctor called. Parts of the CRT display may be underlined or displayed in synchronization with the audio messages.

While a student is on the system he may complete as many as 5 or 10 problems of the type shown above per minute; providing a significant amount of coded lesson material for student use is a major problem. The typical procedure in the reading program is to present material in blocks such that the problems are alike in format, differing only in certain specified ways. The two example problems differ only in 1) film display, 2) words displayed, 3) problem identifier, 4) the three audio numbers, 5) row display of "→" (correct answer row), 6) correct answer area, 7) correct answer identifier. This string of code can be defined once, given a two-letter name, and used later by giving a one-line macro command; the specifics which vary from problem to problem are called parameters.

The use of macros cuts down greatly the effort required to present many different but basically similar problems. For example, the two problems which were presented command by command above would be presented in macro format:

Problem 1: CM PW]F01]bat]bag]rat]A01]ABCD1]A04]A02]A03]7]1,7,3,18]C1]

Problem 2: CM PW]F02]card]cart]hard]]ABCD2]A07]A05]A06]5]1,5,4,18]C2]

The command to call a macro is CM and PW is our two-character code for the macro involving a picture-to-word match. Notice that in problem 2 there is no introductory audio message; the "]]" indicates that this parameter is not to be filled in.

The macro capability of the source language has two distinct advantages over code written command by command. The first is ease and speed of coding. The call of one macro is obviously easier than writing the comparable string of code. The second advantage is increase in accuracy. Not only are coding errors drastically curtailed, but if the original macro is defective or needs to be changed, every occurrence of it in the lesson coding can be corrected by modifying the original macro; in general, the code can stay as it is. The more standard the problem format the more valuable the macro capability becomes. Apart from a few non-standard instructional audio messages and display items, approximately 90-95% of all the reading curriculum has been programmed using roughly 110 basic macros.

The final step in translating the lesson material to a form usable in the computer is the lesson assembly process. A series of machine language programs read in the coded lessons, expand the macros, translate the audio code to actual tape addresses from the audio assembly table and finally read out the assembled lesson translated into binary code onto a disk.

The final editing steps are carried out at the student terminals in a process called debugging. Each lesson is examined in detail by making all correct and all incorrect responses. Errors left uncorrected from any stage in the production process are detected as they might be seen by a student. Corrections are made and the lesson is gone through again in a similar manner. This iterative debugging process is vital for the assurance that the student will not suffer from human error inherent in the complex process of lesson production.

SOME PRELIMINARY RESULTS

Weekly Report.

A report on student progress at a fairly gross level is generated each week primarily for the information of the teachers. The report contains the individual student's name, identification of the lesson on which the student is working, the number of proctor calls received for that student during the week, and a cumulative weighted index of the student's performance for each of the six major problem blocks (i.e., letter discrimination, word presentation, matrix construction, comprehension, compound words, polysyllabic words). Also included in the report are the cumulative number of sessions the student has had on the machine, his number of absences during the week and the total amount of time off the system that week. The last three categories are computed from the teaching proctor records. The remainder of the report is generated by reading pertinent items from the student records. The performance index for each student in each of the six problem blocks is computed in the following manner. A series of six counters are assigned for the computation of the performance index. As a student proceeds through a problem block the counter for that block is indexed to indicate the proportion of initial correct responses made. When the student has completed the block the contents of that register are used to update the associated performance index in the student record. The current value for the register is read out and the new value is added. The updating of the register, however, is done by using a weighting formula: $I_n = I_{n-1} (.40) + I_n (.60)$. The weighting procedure is used to reflect more accurately the student's current performance in relation to his past performance. Computations are done, however, in the integer mode which makes the index somewhat conservative.

Although the weekly performance report is prepared primarily for the classroom teacher and the teaching proctors, it provides sufficient data for examining several rather general questions of overall student performance. The weekly student performance record has been used in the two following analyses.

<u>Spread on Main Line</u>. A central core of problems within the lesson material are considered main line problems in the sense that they are problems which each student must master. A student may be branched around blocks of main line problems by successfully passing certain screening tests. On the other hand, a student may be branched to appropriate remedial material if he has difficulty with these central problems; but in every case he is returned to that set of main line problems for which remedial material was introduced.

Each lesson contains an average of 125 main line problems. Therefore, the number of lessons completed by a student may be used as an index of the number of main line problems successfully completed. Figure 12 shows the number of main line problems completed each week by the fastest, slowest and median student. Figure 12 was derived by identifying the student who had completed the most lessons on the final student progress report for the week of June 14th and, from the same report, identifying the median student and the slowest student, considering only those students who had begun the program on the 15th of November, 1966. New students who had moved into the school and the remedial second-grade students are not considered in the derivation of Figure 12 -- or in any of the data reported in this paper.

Insert Figure 12 about here

The year ended with a difference between the fastest and slowest student of 4,125 problems completed. The inter-quartile range was 1,375 problems and the median student completed 2,625 main line problems. There was, however, a rather wide variation in the amount of time spent on the system by the students. In order to take this variation into account a rate of progress score was computed by dividing the number of problems completed at the end of the year by the number of sessions that the student had on the system. The cumulative rate of progress for the highest, lowest, and median student is shown in Figure 13, expressed in terms of number of main line problems completed per hour of instruction. The range in rate of progress was between 35 problems per hour for the slowest student to 170 main line problems per hour for the fastest student. The inter-quartile range is 45 to 110 with the median at 75 problems per hour.

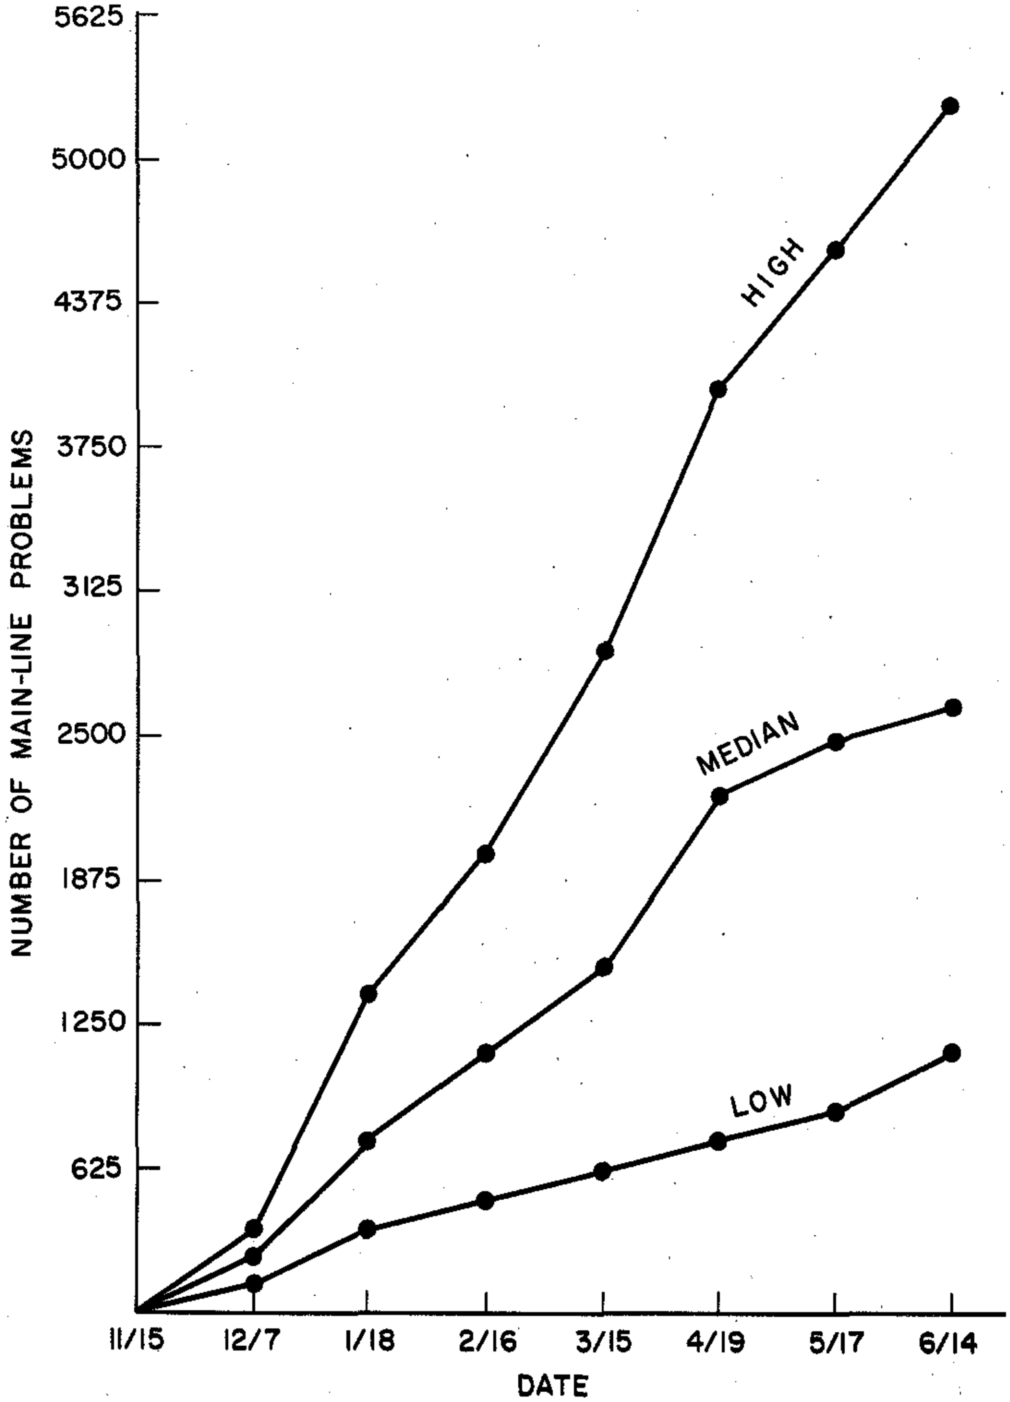

Figure 12.

Cumulative Number of Main-Line Problems Completed

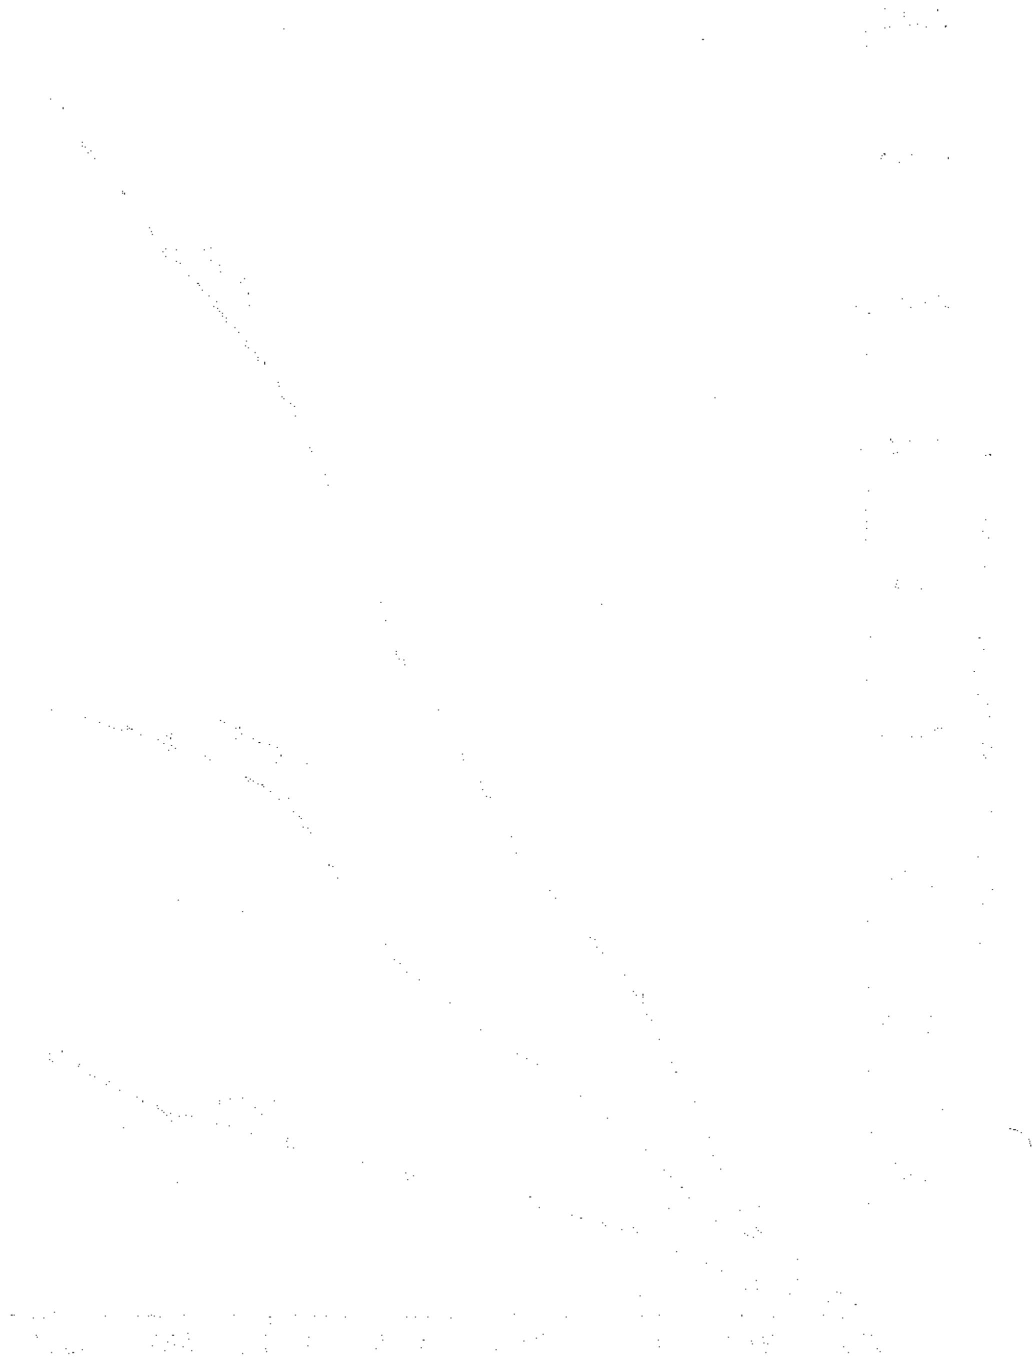

Insert Figure 13 about here

From the standpoint of both the total number of main line problems completed during the year and the rate of progress, it is clear that the CAI reading curriculum is accounting for individual differences on at least one dimension (i.e., the movement of the individual student through the lesson material). The differences noted in Figures 12 and 13 must not be confused with a variation in rate of response. The differences in the rate of response between students was very small; the average rate was approximately four responses per minute. The differences in the total number of main line problems completed and in the rate of progress can be accounted for by the amount of remedial material, the optimization routines, the number of corrections, and the number of accelerations for the different students.

Sex Differences. It has been a common finding in reading studies that girls generally surpass boys in the acquisition of reading skills and in reading performance, particularly in the primary grades (Gates, 1961; Wyatt, 1966). The suggestion has been made many times that these differences might be attributed, at least in part, to the social organization of the classroom and to the value and reward structures of the predominantly female primary-grade teachers. It has also been argued that because of differences in developmental rates first-grade girls are more facile in visual memorization than boys of the same age -- a facility which would favor the girls in the essentially sight-word method of vocabulary acquisition commonly used in the current basal reading series. If these two arguments are viable, then one would expect that placing students in an asocial environment such as a CAI tutorial system and presenting a linguistically oriented curriculum which emphasizes analytic skills, as opposed to rote memorization of words, would minimize the sex difference in reading performance. In order to test this notion, the rate-of-progress scores taken from the final teachers' report were rank ordered and tested for significant sex effects using a Mann-Whitney U-Test. The null hypothesis in this and in the following tests is that the scores for the boys and the scores for the girls have the same distribution. The test of sex effects yielded z of .05. Under the null hypothesis the probability of z being greater than or equal to 0.35 is 0.36. Sex difference then is not an influential variable in the rate of progress in the Stanford CAI reading curriculum.

Sex differences still might be an influential factor in accuracy of performance. To test this notion the final performance index scores for each of the four standard problem blocks reported on the weekly teachers' report were rank ordered and examined under the Mann-Whitney U-Test. The results were as follows:

Letter identification; $Pr(z \geq 0.33) = 0.37$.

Word list; $Pr(z \geq 1.83) = 0.03$.

Sounding matrix; $Pr(z \geq 1.41) = 0.08$.

Sentence comprehension; $Pr(z \geq 1.37) = 0.09$.

The only significant difference, at the 0.05 level, was found in the word list scores. However, the scores in the matrix and comprehension sections were in the expected direction (i.e., girls excelling boys). These results, while by no means definitive, do lend support to the notion that when students are removed from the normal classroom social milieu and placed in the asocial

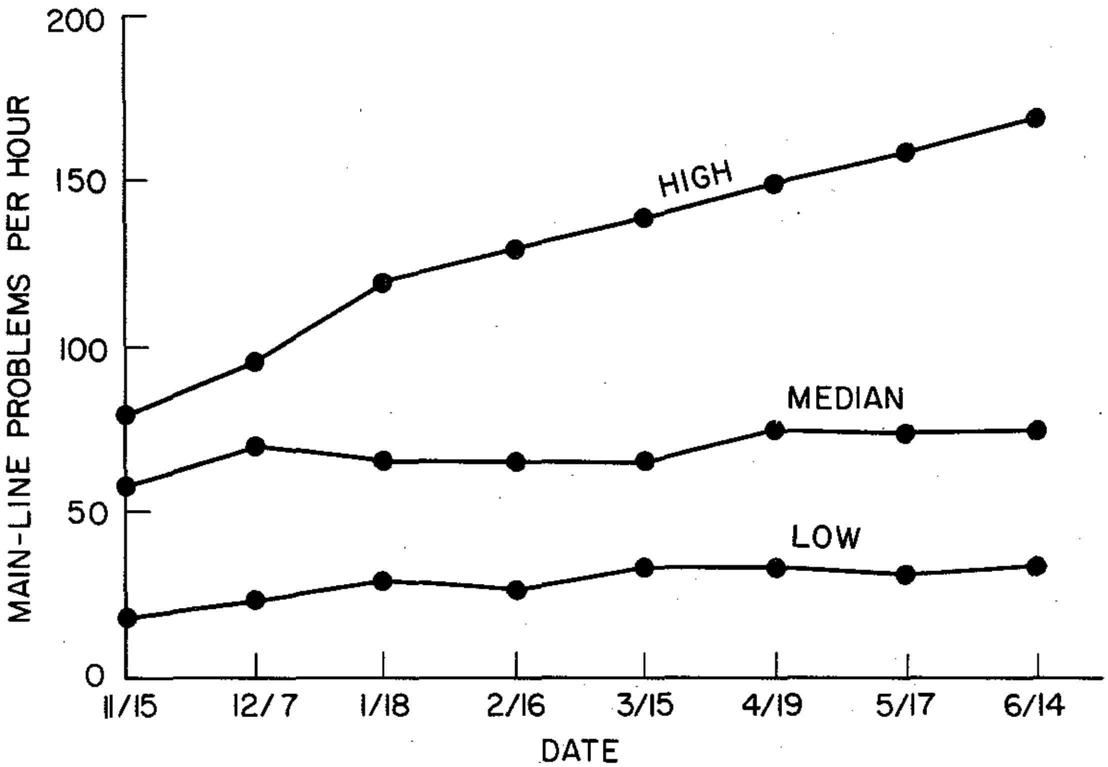

Figure 13.

Cumulative Rate of Progress of Fastest, Median, and Slowest Students

environment of a CAI tutorial system, boys perform as well as girls in overall rate of progress through a reading curriculum. The results also tend to support the idea that in a CAI environment the sex difference is minimized in direct proportion to the emphasis on analysis versus rote memorization in the learning task. The one problem section where the girls achieved significantly higher scores than the boys was the word list section which is a paired-associate list learning task.

Analysis of the Response Data.

Each student response is first stored in an area of core memory called a response buffer. When the limit of the storage capacity of the response buffer is reached, the contents of the buffer are read out on a response tape. The steps in the preparation of the response tape for storage and for serious analysis of the student's response data is illustrated in Figure 14. A compressed response tape is generated from the data response tapes by means of a compressor routine which removes the inter-record gaps on the daily response tape. It also may be used to compress two or more previously compressed response tapes. It is the compressed response tape which is held in permanent storage.

Insert Figure 14 about here

There are items in the student response record which are of less than immediate interest insofar as data analysis is concerned. Therefore, the compressed response tape is subjected to a strip routine which strips out those items which will be used in immediate analysis to form a subset response tape. In the case of the reading program the items stripped out are the student number, course name, response identifier, the match identifier, and the latency record. The strip routine number 1 may also be used to select certain records of interest from the compressed response tape.

A generalized sort/merge program is applied to the subset response tape which will handle up to ten key words in any order within a record and which preserves chronological order in cases where two or more entries are the same. By means of the generalized sort/merge routine a sorted subset tape is generated where the records are recorded first according to student number and second according to problem number (response identifier). A tape-to-printer routine is associated with each tape, from the daily response tapes to the sorted subset tape. This routine allows for listing of the particular tape which then may be examined for completeness and accuracy. It is the sorted subset tapes which are used as input for data analysis programs. The first data analysis program which will be discussed here is in reality a data description program at the level of the problem type.

Due to the complexity of the branching routines within a problem type and the variation of the branching logic between problem types, it is possible to have an indeterminate number of responses to a given problem by a given student. These responses then must be categorized. A reasonable method of categorizing such responses is illustrated in Figure 15.

Insert Figure 15 about here

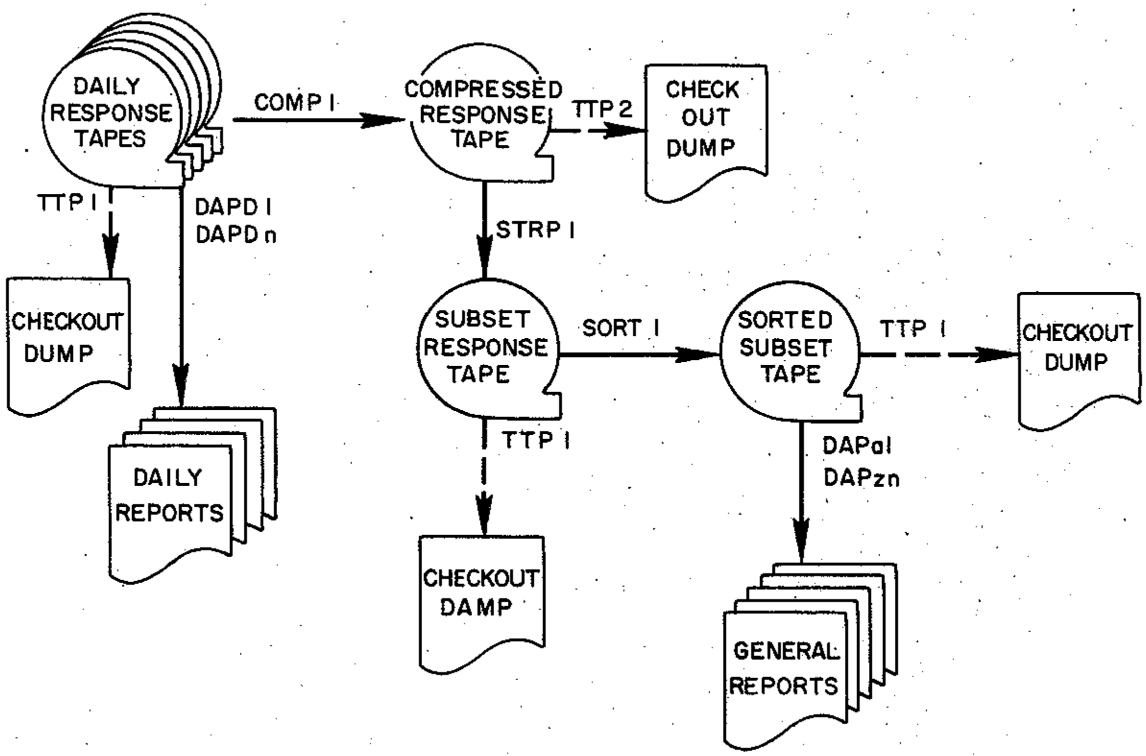

COMP1: COMPressor routine number 1.

DAPxn: Data Analysis Program x number n.

SORT1: SORT routine number 1.

STRP1: STRiP routine number 1.

TTP1: Tape To Printer routine number 1.

TTP2: Tape To Printer routine number 2.

Figure 14.

Data Analysis Programs

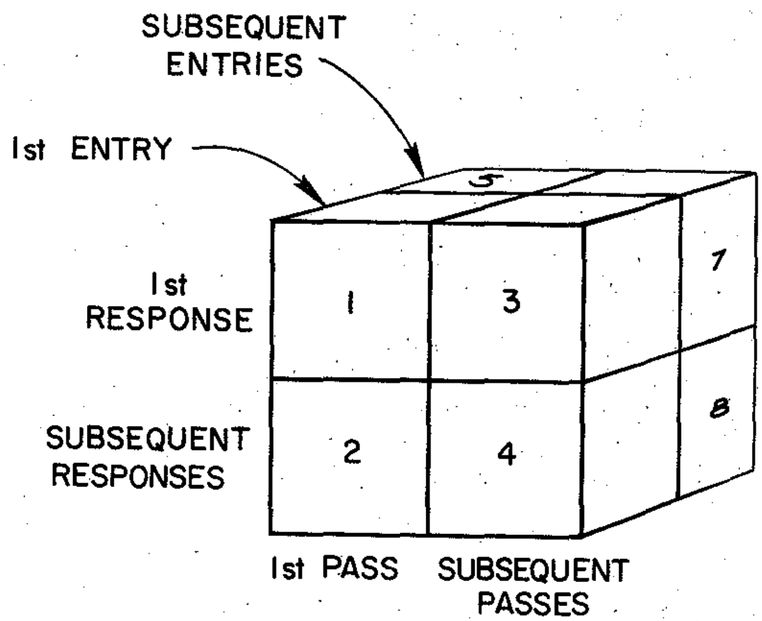

Figure 15.

Response Categorization at the Problem-Type Level

The initial response to a problem will fall into cell 1. If the response is incorrect, the student will be required to respond again to the same problem. Subsequent responses, until a correct response is made, will then fall into cell 2. If the problem type uses an optimization routine, then the list of problems will be gone through again, deleting those problems which received an initial correct response. On each optimization run there will be a first time for a student to respond to each problem. Those first responses will fall in cell 3. However, as in cell 1, that first response may be in error, in which case subsequent responses will be called for on that problem which will then fall in cell 4. There are many cases in the curriculum where a student will be branched out of a problem type into some appropriate remedial material and then be returned to the original problem type. The initial response to a problem in a subsequent entry to a problem type will fall in cell 5 and corrections in cell 6. The initial responses of optimization runs on subsequent entries will fall in cell 7 and subsequent correction responses will fall in cell 8.

This method of classification of responses uniquely defines each response by cell number. Given a sorted subset tape, as specified above, a general algorithm was developed to assign a cell number to each response. The general algorithm is as follows: For a given student and a given problem number the first response encountered is assigned to cell 1. The second to n^{th} responses are assigned to cell 2 where n equals the first CA match identifier. The $n+1$ response is then assigned to cell 3. The following responses with WA match identifiers, if any, and the next CA match identifier are assigned to cell 4. The next response is then assigned to cell 5. The remaining instructions are essentially a repetition of the first portion of the algorithm. It must be emphasized, however, that this is a general algorithm to which several variations must be added to account for the unique branching logic of many of the problem types.

In addition to the above described cell identification each response is also tagged with a restart run number. This restart identification is required since a student who is signed off for any reason (e.g., end of the period, machine failure) while he is at some point within a problem block can only start again at the previous restart point. This situation implies that the probability is high that at the beginning of a day a student will be repeating some indeterminate number of problems from the day before. Those responses which result from repeated exposure to problems because of restart runs, must be treated in data analysis in a different manner than those responses which were given to the original exposures to the problems. The association of problems to restart runs is done during the sort routine between the subset response tape and the sorted subset tape.

With each student response tagged with a cell number and a restart number, the process of compiling a series of summary reports was begun. Several levels of summary reports can be generated. Reports can be generated which describe the performance of the total class of students at the level of the problem block, the problem type, and the individual problem. Other reports can be generated which will summarize performance of individual students at the three levels: block, type, or problem. The first series of reports which are in the process of generation at the time of this writing, describe the performance of the total student group at the level of the problem type. A separate report is produced for each problem type within each lesson. The report describes the mean, standard deviation, and range of the distribution within each of the eight cells for: the number of problems, number of responses, the proportion of responses which were incorrect, the proportion of problems with at least one error, the correct answer latencies, the wrong answer latencies, the overall latencies and the CA, WA, and overall speed scores.

Achievement Tests.

Even though this first year's operation of the system has been viewed as essentially an extended debugging process, it was felt that a comprehensive end-of-the-year testing program might yield some interesting insights to the potential impact of the program on overall reading achievement. Accordingly, a battery of tests was assembled which would measure achievement in each of the major areas of reading behaviors taught at the first-grade level. The behaviors identified and included in the testing program were:

1. Identification of letter and letter-strings.
2. Acquisition of an initial sight vocabulary.
3. Acquisition of word decoding skills:
    a. initial consonants
    b. initial consonant clusters
    c. medial short vowels
    d. final consonants
4. Acquisition of reading comprehension skills:
    a. word meaning
    b. syntax or form class for sentence comprehension
    c. paragraph comprehension:
        1. main idea
        2. recall of facts

Wherever possible we chose standardized tests of the reading behavior to be evaluated. In several cases, the tests only approximated what we were trying to measure or were not pure tests of a single reading behavior. As far as we could determine, none of the available test batteries included tests of all the reading behaviors listed above. It was particularly difficult to find group administered tests measuring a child's ability to decode words and the child's knowledge of syntax or form class behaviors, therefore Ruth Hartley of the project staff devised tests of these behaviors.

The tests which were chosen for this test battery were derived from the following sources:

1. Gates-MacGinitie Reading Test, Primary A, Form 1, 1965 edition.
2. California Reading Test, Lower Primary, Form W, 1957 edition with 1963 norms.
3. Stanford Achievement Test, Reading Sections, Primary 1 Battery, 1964 edition.
4. Project developed tests.

The students in the CAI mathematics program provided an ideal control group. Analysis of the individually administered Stanford Binet I.Q. tests at the beginning of the school year indicated that the students in the reading program and the students in the mathematics program may be considered as two samples from the same population. The mean I.Q. for the reading group was 92.5 (standard deviation 15.6), and the mean I.Q. for the mathematics group was 91.8 (standard deviation 14.6). Any "Hawthorne effect" which might be induced by the CAI experience is controlled since the mathematics students had an equal amount of time on the system but for a different subject matter. The mathematics students received a program of reading instruction which might be termed traditional, relying primarily upon the Ginn and the Alyne and Bacon first-grade readers.

All tests, with the exception of the Stanford Achievement Test, were administered by members of the project staff, assisted by graduate students from the School of Education at Stanford University. Each testing team, consisting of a project staff member and a graduate student, administered a given test for all of the students in order to eliminate tester effects. The Stanford Achievement Test was administered by the individual classroom teachers as part of the testing program of the State of California. All tests were conducted within the normal classroom setting. The testing program was carried out during the month of May, 1967.

The content and characteristics of the three standardized tests in the battery are well-known and need not be reiterated here, however, a brief description of the project developed tests might be helpful in the interpretation of the results.

Vocabulary Test. This test consisted of words and patterns that were presented in the CAI program. Some of the words and all of the word patterns included in this test also appear in the Ginn and Alyne and Bacon first-grade basic readers. The test consisted of 20 sets of three words each. The student was directed to circle the word pronounced by the examiner.

Phonetic Discrimination. This test was designed to check the child's ability to distinguish between minimal contrasts of monosyllabic words and to connect the auditory sound of a phoneme with its grapheme. The test contained two sections: 1) the student was required to identify in a set of three words the word pronounced by the test administrator (e.g., bat, rat, sat) and 2) the student was required to fill in the missing letter or letters in a word pronounced by the administrator.

In both cases the test dealt with initial consonants, initial consonant clusters, medial short vowels and final consonants.

Reading Comprehension: Understanding of Syntax or Form Class Behaviors. The twenty items of this test consisted of a series of sentences in which one word was omitted. The subject was required to choose, from a set of four words, the word which would complete the sentence. The distractors in the multiple-choice set consisted of one word which was the correct form class but was semantically irrelevant and words from two incorrect form classes. The vocabulary used in this test was derived from the first forty lessons of the CAI program and the Ginn and Alyne and Bacon first-grade readers.

Pronunciation. The pronunciation test was composed of forty items, 20 of which were real monosyllabic words and 20 items were pronounceable nonsense monosyllabic units. Each nonsense item was composed as a minimal contrast to a real word item, (e.g., man,fam). Five pairs were constructed for each of four categories of minimal contrast: initial consonant, medial vowel, final consonant and initial consonant blend. Each item was typed in lower case primary type on a 3 x 5 card and presented to the subjects one at a time. The subjects were required to attempt a pronunciation of each item.

Student responses were recorded on magnetic tape for future more detailed analysis. The results reported in this paper are those of the all-or-none scoring procedure (i.e., an item was scored as correct only if all phonemes of the item were correctly pronounced and in proper sequence).

Overall Results. The results of each of the above tests and their major subsections were examined in a series of three-way analyses of variance (treatment, high/low I.Q., sex). No significant interactions were found in any of the analyses. The expected significant differences on the I.Q. and sex variables were found throughout the tests and are of little interest here. The means and standard deviations for each test for the experimental and control groups are shown in Table 4. It is interesting to note that with the exception of two cases (Stanford Achievement Test, paragraph reading and total score) the direction of differences between the means is in favor of the experimental group.

Insert Table 4 about here

The test battery may be divided into two major categories: 1) tests designed to evaluate the goals and linguistic orientation of the CAI program, and 2) tests designed to evaluate the outcomes of a quite different approach to initial reading (i.e., the traditional basal reading series approach). Differences between the means of each of the project developed tests were statistically significant and in one case, the pronunciation test, the differences were fairly dramatic. Statistical significance was also achieved in three of the eight sub-scales of the standardized tests. Two of those sub-scales, however, the vocabulary section of the California Achievement Test and the word-study skills section of the Stanford Achievement Test, were composed of a series of tests of sub-skills. It was necessary, therefore, to investigate the possibility that large differences on one or two of the sub-skills might be exerting an influence powerful enough to produce significance in the sub-scale as a whole. Accordingly, the scores on each test of sub-skills were examined in the same three-way analysis of variance as described above. The results may be seen in Table 5. The significance levels held in all four of the sub-skills for the word-study section of the Stanford Achievement Test and in three of the five tests of sub-skills in the vocabulary section of the California Achievement Test.

Insert Table 5 about here

The impact of this year's CAI instruction on reading performance may be more readily visualized when the results are tabulated by reading behaviors which were measured in the evaluation program as in Table 6. Eight of the nine tests of decoding skills resulted in differences in favor of the CAI group which were significant at the .05 level and in six of those eight tests the significance was at the .01 level or beyond. Three of the tests of comprehension at less than the paragraph level resulted in differences in favor of the CAI group which were significant at the .05 level. No significant differences were found in the tests of comprehension at the paragraph level.

Insert Table 6 about here

The above results are most encouraging for the potential impact of CAI on initial reading in light of the fact that even the fastest student in this year's program progressed only a small fraction of the way through the first year's curriculum (Figure 16). The curriculum was designed with the expectation that the able student would be able to complete approximately 180 lessons in the first year. As may be seen in Figure 16, the top student in this year's program completed only approximately 22 percent of the expected total number of lessons. The average student in this year's group completed only 11 percent of the total first year's program.

Insert Figure 16 about here

A comparison of Table 1 and Figure 16 will verify that no student in this year's run had been exposed on the CAI system to more than five basic patterns:

Table 4.

MEANS AND STANDARD DEVIATIONS: MAJOR TEST SECTIONS

	Experiment		Control		Significance Level
	Mean	S.D.	Mean	S.D.	
Gates*					
Vocabulary	41.1	9.26	39.3	9.01	N.S.
Comprehension	39.5	8.93	38.8	9.12	N.S.
Total	38.5	9.07	37.4	9.02	N.S.
C.A.T.*					
Vocabulary	45.9	1.76	38.1	2.02	$p < 0.01$
Comprehension	41.4	2.51	40.6	2.16	N.S.
Total	45.6	0.58	39.6	1.01	$p < 0.01$
S.A.T.*					
Word Reading	40.3	4.84	40.0	2.44	N.S.
Para. Reading	36.0	9.73	38.5	4.07	N.S.
Vocabulary	41.7	3.89	40.7	2.05	$p < 0.01$
Word Study	45.9	3.72	44.9	2.50	$p < 0.01$
Total	44.2	5.89	44.6	3.76	N.S.
Project**					
Form Class	8.6	5.23	6.4	4.95	$0.01 < p < 0.05$
Vocabulary	17.7	4.16	15.7	4.27	$p < 0.01$
Pronunciation	11.28	10.21	4.94	7.66	$p < 0.01$
Phonetic Discrim.	14.06	20.07	10.63	15.32	$p < 0.01$

*Standard Scores: M = 50, S.D. = 10

**Raw Scores

Table 5.

MEANS AND STANDARD DEVIATIONS: SUB-SCALES

	Experiment		Control		Significance Level
	Mean	S.D.	Mean	S.D.	
C.A.T.					
Vocabulary:					
Word Form	18.3	5.13	15.0	5.30	$p < 0.01$
Word Recognition	13.4	4.87	11.7	4.50	$0.01 < p < 0.05$
Opposites	6.2	2.96	5.7	2.70	N.S.
Picture Association	8.2	2.92	6.8	2.26	$p < 0.01$
Letter Recognition	21.2	5.17	20.6	6.20	N.S.
S.A.T.					
Word Study:					
Audio Perception					
Beginning Sounds	9.87	2.71	8.00	2.70	$p < 0.01$
Ending Sounds	8.89	2.77	6.78	3.13	$p < 0.01$
Phonics	10.20	2.54	8.89	2.70	$p < 0.01$
Phonograms	7.28	2.88	6.33	2.18	$0.01 < p < 0.05$

Table 6.

SIGNIFICANCE LEVELS OF TESTS OF READING BEHAVIORS

Reading Behaviors Measured	Test*	Significance Level .01	.05	N.S.
A. Decoding				
1. Letter and Letter String Identification	1			x
	2	x		
2. Initial Sight Vocabulary	3		x	
	4[p]	x		
3. Word Decoding Skills	5	x		
	6	x		
	7		x	
	8	x		
	9[p]	x		
	10[p]	x		
B. Comprehension				
1. Word Meaning from Picture Cues	11			x
	12			x
	13	x		
2. Word Meaning - Vocabulary	14	x		
	15			x
3. Understanding of Syntax and Form Class	16[p]		x	
4. Sentence, Paragraph and Story Comprehension	17			x
	18			x
	19			x

Table 6 (cont'd)

	*Test	Number of Items
1.	C.A.T., Letter Recognition Test	24
2.	C.A.T., Reading Vocabulary-Word Form Test I, Section A	25
3.	C.A.T., Reading Vocabulary-Word Recognition, Test I, Section B	20
4.	Hartley Vocabulary Test	20
5.	S.A.T., Word Study Skills Auditory Perception of Beginning Sounds	14
6.	S.A.T., Word Study Skills Auditory Perception of Ending Sounds	14
7.	S.A.T., Word Study Skills Phonics - Recognition of Word from Spoken Cues	14
8.	S.A.T., Word Study Skills Phonograms - Rhyming Words	14
9.	Hartley, Phonetic Discrimination	40
10.	Hartley, Pronunciation Test	20
11.	Gates, Vocabulary Test - Picture Association	48
12.	S.A.T., Word Reading - Picture Association	35
13.	C.A.T., Vocabulary - Picture Association Test I, Section D	15
14.	S.A.T., Vocabulary	39
15.	C.A.T., Vocabulary - Meaning of Opposites, Test I, Section C	15
16.	Hartley: Form Class	20
17.	Gates, Comprehension Test	34
18.	C.A.T., Comprehension - Test 2	15
19.	S.A.T., Paragraph Meaning	38

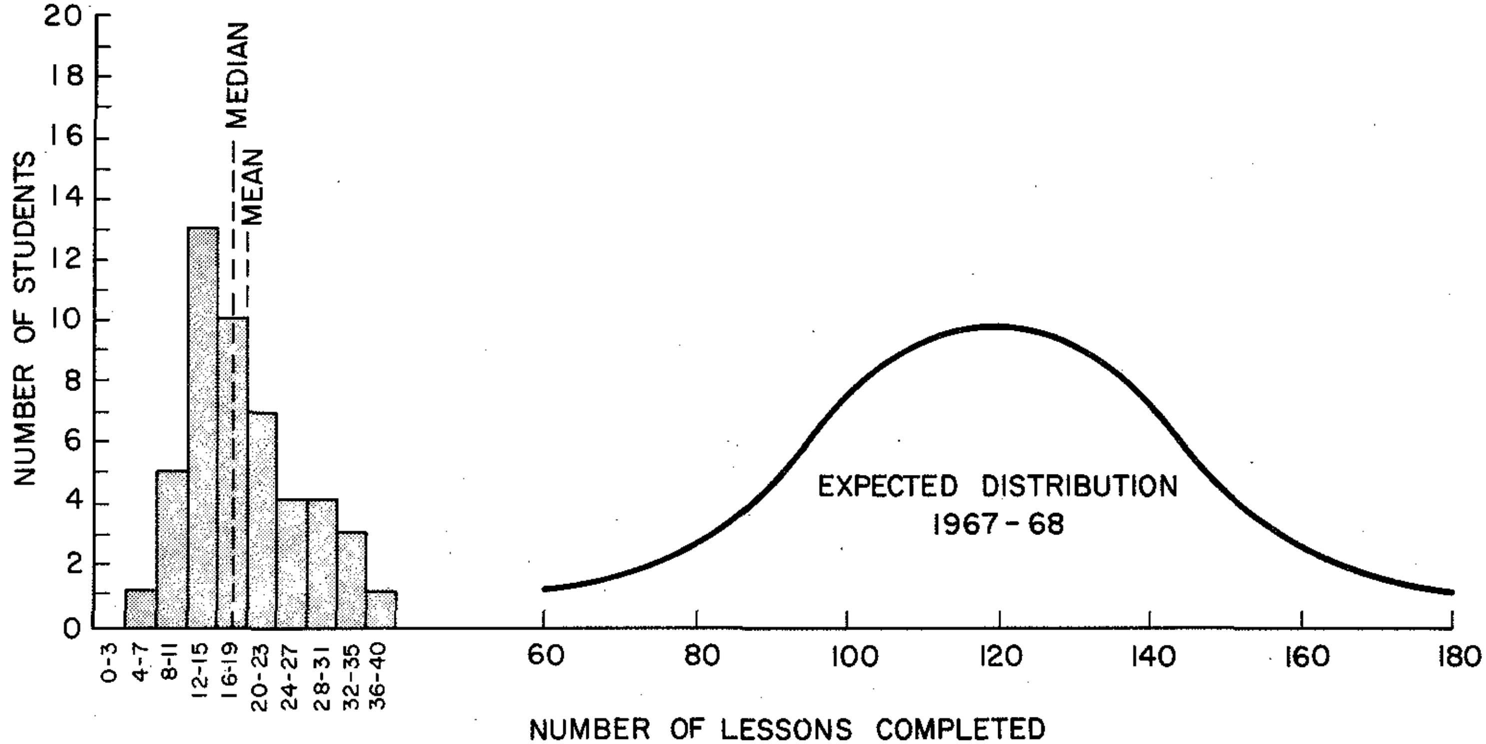

Figure 16.

Distribution of Students Over Lessons (1966-67)

ac, cac, ccac, ic, and cic (where "c" stands for any consonant and each vowel is specified). This minimal exposure was sufficient to produce a significant increase in performance in nine out of the ten tests of decoding skills. Some transfer of learning is evidenced by the fact that six of those nine tests were commercial standardized tests whose vocabulary and word patterns were not based on the Stanford CAI program. Transfer of training is also born out by the observation of the classroom teachers involved, particularly the teacher of the low maturity group who has reported that a noticeable percentage of the low group of students were discovering word patterns in their classroom reading instruction and generalizing those word patterns in the attack of new words.

The most consistent results were found in the area of decoding skills which is exactly the area of initial reading which is handled in greatest detail in the Stanford CAI program. On the other hand, results in the comprehension tests were mixed, with no significant differences being found in tests of paragraph comprehension. It must be noted that even the fastest student in this year's program had not reached the level in the lesson materials where exercises on the comprehension of connected discourse is introduced. It is also true, as stated earlier in this paper, that the comprehension exercises in the Stanford CAI curriculum are more of the nature of data gathering devices than they are a well defined and complex teaching effort.

FINAL REMARKS

It has been our goal in this paper to describe the Stanford Reading Project in terms of the CAI system, its operation, its function in the elementary school setting, the rationale and major components of the curriculum, and to give some feeling for the processes involved in the production of such a curriculum. We have also discussed some results from the first year operation of the program. It cannot be emphasized too strongly, however, that the data and the conclusions drawn from them are at best tentative; they are based on the first twenty-six weeks of operation which must be viewed essentially as an extended debugging period for both the system and the learning materials. The progress of the students was limited by known problems which are in the process of being remedied for the next year's operation. Revisions for next year involve modification of: the branching structure, the display formats, and the audio search procedures. The projected distribution of students at the end of the school year 1967-68, shown in Figure 16, is based upon an experiment run during the last week of the 1966-67 operation in which all of the students were placed in selected lessons of advanced material where the known problems had been corrected. The projection is based not on a pious hope but on some quantity of empirical data. The results of this first year's evaluation, while tentative, do suggest that such a program undertaken at the end of the 1967-68 school year may demonstrate an impact of a CAI program on the acquisition of initial reading skills which is not only statistically significant but practically significant as well.

As the curriculum becomes refined and the system achieves a high degree of stability and reliability, the Stanford-Brentwood CAI laboratory will be made available as a laboratory for educational research in the classical sense. It will be relatively simple, for example, to design and implement alternative teaching strategies which may be evaluated in a highly controlled environment.

The problems of optimal teaching strategies may be investigated in a rigorous manner. It is hoped that basic educational research, which may be carried out in such a facility, will prove relevant not only to computer-assisted instruction but will be generalizable also to the important problems of teaching and learning throughout our educational system.

References

Atkinson, R. C. Instruction in initial reading under computer control: the Stanford Project. Technical Report 113, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1967. (To be published in the Journal of Educational Data Processing, Vol. 4, No. 4, Fall, 1967.)

Atkinson, R. C., G. H. Bower, and E. J. Crothers. An Introduction to Mathematical Learning Theory. New York: John Wiley and Sons, Inc., 1965.

Atkinson, R. C. and R. M. Shiffrin. Human memory: a proposed system and its control processes. Tebhnical Report 110, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1967. (To be published in The Psychology of Learning and Motivation: Advances in Research and Theory (Vol. II), edited by K. W. Spence and J. T. Spence, New York, Academic Press.)

Bloomfield, L. Linguistics and reading. Elementary English Review, 1942, 19, 125-130.

Bloomfield, L. and C. L. Barnhart. Let's Read: A Linguistic Approach. Detroit: Wayne State University Press, 1961.

Carterette, E. C. and M. H. Jones. Phoneme and letter patterns in children's language. In Proceedings of Symposium on the Psycholinguistic Nature of the Reading Process. Wayne State University, 1965.

Fries, C. Linguistics and Reading. New York: Holt, 1963.

Gates, A. I. Sex differences in reading ability. The Elementary School Journal, Vol. 61, May, 1961, 431-434.

Groen, G. J. and R. C. Atkinson. Models for optimizing the learning process, Psych. Bull., 66, 4, 1966, 309-320.

References (cont'd)

Hansen, D. N. and T. S. Rodgers. An exploration of psycholinguistic units in initial reading. Technical Report 74, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1965.

Rodgers, T. S. Linguistic considerations in the design of the Stanford computer-based curriculum in initial reading. Technical Report 111, Institute for Mathematical Studies in the Social Sciences, Stanford University, 1967.

Thorndike, E. L. and C. L. Barnhart (Eds.), Thorndike Barnhart Beginning Dictionary. New York: Scott, Foresman and Co., 1964.

Thorndike, E. L. and I. Lorge. The Teacher's Word Book of 30,000 Words. New York: Columbia University, Bureau of Publications, 1944.

Wright, W. W. The Rainbow Dictionary. Cleveland, Ohio: The World Publishing Co., 1959.

Wyatt, Nita M. Sex differences in reading achievement. Elementary English, Vol. 43, No. 6, Oct. 1966, 596-599.

CAI Programming Systems User's Guide: Preliminary Draft. IBM, Armonk, New York, 1967.