

**Best
Available
Copy**

AD-766 459

IMPROVING INTERACTIVE CAPABILITIES IN COMPUTER-
ASSISTED INSTRUCTION

BOLT BERANEK AND NEWMAN, INC.

PREPARED FOR
ADVANCED RESEARCH PROJECTS AGENCY

AUGUST 1973

DISTRIBUTED BY:

NTIS

**National Technical Information Service
U. S. DEPARTMENT OF COMMERCE**

BOLT BERANEK AND NEWMAN INC
CONSULTING • DEVELOPMENT • RESEARCH

BBN Report No. 2631
Job No. 11548, 11794

AD 766459

IMPROVING INTERACTIVE CAPABILITIES IN
COMPUTER-ASSISTED INSTRUCTION

Semi-Annual Technical Report for six months ending July 31, 1973

August 31, 1973

Allan M. Collins, Principal Investigator (617/491-1850)
Joseph J. Passafiume
Laura Gould
Jaime G. Carbonell

DDC
RECEIVED
SEP 11 1973
B

Contract No. N00014-71-C-0228 dated March 1, 1971
Amendment Modification No. P00002, dtd January 1, 1973
Expiration Date, December 31, 1973
Total Amount of Contract, \$358,632.00

Sponsored by: Office of Naval Research
Contract Authority No. NR 154-330
Scientific Officers: Dr. Marshall Farr and
Dr. Joseph Young

and

Advanced Research Projects Agency
ARPA Order No. 2284 dtd 30 August 1972
Program Code No. 61101E

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency, the Office of Naval Research, or the U.S. Government.

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
US Department of Commerce
Springfield, VA. 22151

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

TABLE OF CONTENTS

	Page
1. SUMMARY	1
2. PRESENTATION MODES IN SCHOLAR	4
2.1 Introduction	4
2.2 Descriptions of Both Modes	6
2.3 Testing	14
2.4 Results and Discussion	18
3. TEACHING PROCEDURAL KNOWLEDGE	22
3.1 Introduction	22
3.2 Preliminary Analysis of Dialogues on the ARPA Network	22
3.3 Choice of the TNL5 System	24
3.4 Aspects of Teaching Procedural Knowledge	28
3.5 Future Plans	35
4. TEACHING GEOGRAPHY WITH MAPS	37
4.1 Introduction	37
4.2 Topic and Question Generation	37
4.3 Answer Evaluation and Error Correction	39
5. REFERENCES	43

SECTION 1

SUMMARY

Our work on the development of interactive capabilities in the SCHOLAR CAI system during the last six months centered in three main areas: (1) implementation of two presentation strategies in SCHOLAR (Tutorial mode and Block-Test mode) and a comparative evaluation of these two modes using high school students as subjects; (2) initial study based on analysis of tutorial dialogues of how to teach procedural knowledge interactively within SCHOLAR, and (3) addition of a module for teaching geography using the map display and related question-answering facilities recently added to SCHOLAR. Each of these three areas comprises one section of the following report.

The work in the first area involved development of two large modules for the SCHOLAR system. Initially SCHOLAR (Carbonell, 1971) did not present material except to answer questions. Both new modules select topics to be discussed and then present material and ask questions about the topics selected. Tutorial mode is based on extensive analysis of dialogues between different tutors and students, performed earlier under this contract. In this mode SCHOLAR first questions the student to find out what he knows about each topic, and then presents some related information limited to what the student can assimilate. Block-Test mode is based on the strategy used in programmed instruction. In this mode SCHOLAR first presents information and then asks questions about the information presented.

When these modules were completed we ran a small experimental study with eight high school students to compare the two modes.

Each student learned about two South American countries in one mode and two countries in the other mode. (We only permitted them to ask questions to SCHOLAR in Tutorial mode.) The amount of learning in each mode was measured by the difference in test scores on a pre-test and post-test given a couple of days before and after the teaching sessions. The results indicated a significant difference in favor of the Tutorial mode. We plan to make improvements in both modes along lines suggested by the students and carry out further testing to explore systematically what are effective teaching strategies.

In the second area we have been conducting tutoring sessions where tutors interactively teach students with varying backgrounds how to use a computer system. Then we analyze these tutorial dialogues using protocol analysis in order to determine what strategies are effective for teaching procedural knowledge. The most salient fact that emerged from the initial analysis was the necessity for the student to try out what he learns as he learns it. This led us to the decision to attempt to embed a version of the system being taught within SCHOLAR so that the student could interact with SCHOLAR while trying out what he learns. Two other aspects of teaching procedural knowledge that emerged from the dialogue analysis was the importance of explaining procedures both in general terms and with respect to the specific example at hand, and the usefulness of explaining new procedures in terms of their similarities and differences with known procedures. We are now starting to develop new modules for the SCHOLAR system that embody these and other ideas derived from the dialogue analysis.

In the third area we have started to develop a teaching module to utilize the map display in SCHOLAR (Warnock & Collins, 1973). Because the visual representation of the maps in SCHOLAR is

highly integrated with the semantic network of facts about South America, the student can control the display verbally. The new module will allow SCHOLAR to ask the student to locate different places by pointing, and to name and point to specific places, such as the major cities in Argentina. When this module is completed, we plan to integrate it with the teaching modes described above and use it in further testing with high-school students.

SECTION 2

PRESENTATION MODES IN SCHOLAR

2.1 INTRODUCTION

To produce computer environments which result in truly individualized learning, the computer must generate material and questions based on its knowledge of the subject matter and the user. In order to see what is involved in individualized instruction, we studied in a previous contract how the human tutor adapts his teaching to the individual student. To do this we made an in-depth analysis of dialogues on South American geography between human tutors and students (Collins, Carbonell, & Warnock, 1973). Using the concept of subroutines, we analyzed the tutor's behavior with regard to error-correction, question generation, dynamic generation and handling of an agenda, and selection of most relevant material in presentation or in answers.

To investigate the effectiveness of tutorial instruction we implemented two strategies in the SCHOLAR CAI system. The two strategies are called Tutorial mode and Block-Test mode. The former is based on the tutorial dialogue analysis. The latter is a variation of the presentation strategy used in traditional CAI systems. Under both strategies, information is covered exhaustively in the order of importance, as measured by I tags in the data base. Once a topic has been selected, such as location, all of the information under that topic is discussed down to a pre-specified but adjustable level. The level is adjusted during the dialogue, depending on how much time is left.

In Block-Test mode, SCHOLAR first presents a paragraph of information, then it questions the student about the information

just presented. Errors are corrected by providing the correct answer.

In Tutorial mode, SCHOLAR starts by asking a question rather than by presenting material. SCHOLAR goes deeper into the topic, down to the prespecified level for as long as the student can answer correctly. If the student cannot answer, SCHOLAR gives the correct answer, explains any incorrect answer, and provides some related pieces of information about the correct answer. SCHOLAR then goes on questioning by backing up one level in the network. The questions are mostly WH-questions and fill-in-the-blank type questions, except for some true-false type questions to avoid open-ended answers by the student. In this mode, after the material is covered at a fairly shallow level on a first pass, SCHOLAR starts reviewing, using the same basic strategy. When a question was answered earlier or the information was presented to the student, SCHOLAR asks about it when reviewing. In addition, new information is presented, which is related to the old. Reviewing continues until the time is used up. Lastly, in this mode, the student is allowed to ask questions of the form, "What is X?", "Where is X?", or "Tell me about X.", in order to ask about any unfamiliar terms or concepts that might be used by SCHOLAR.

Over the past few months, teaching sessions with eight high-school students were conducted. Students learned about two countries in one mode, and two in another mode. This way, the student served as their own controls in comparing the learning as reflected in the pre- and post-test scores. At the same time, students were questioned about which aspects of the different modes they found helpful and which aspects not helpful. The results showed that the students learned significantly more in the Tutorial mode than in the Block-Test mode.

In the succeeding sections we shall describe in greater detail how the two modes were implemented, how the study was conducted, and the results of that study.

2.2 DESCRIPTIONS OF BOTH MODES

We shall describe how the Tutorial mode and the Block-Test mode work by explaining how a portion of the dialogue with one of the high-school students was generated.

Tutor-SCHOLAR starts by choosing a major context, usually specified on the preset agenda. SCHOLAR begins by asking, "What is Argentina?" (see Fig. 1). This question is generated by entering the node Argentina (see Fig. 2) which was specified on the agenda, and choosing the first most important (as measured by I tags) unused topic within a prespecified, but adjustable level. The level is the maximum allowable semantic depth, measured from Argentina, of the information that will be discussed. It is adjusted according to the time remaining. Under the node Argentina, the first most important topic is the superordinate (SUPER). Selecting this element, SCHOLAR constructs the question, "What is Argentina?". The student's response is "county," but SCHOLAR assumes that the student intended to say "country," points out the misspelling, and accepts the answer. At this time, a hash tag (see Appendix) is written, which denotes the nature of the interaction and refers to the element of the data base (i.e., SUPER NIL COUNTRY). Whether a piece of information will be discussed is always a function of how the information was used in the prior discussion.

WHAT IS ARGENTINA?

COUNTY

YOU SHOULD HAVE WRITTEN "COUNTRY".
VERY GOOD.

IS IT TRUE THAT ARGENTINA IS IN SOUTHERN SOUTH AMERICA?

T

YOU ARE CORRECT.

THE EASTERN COAST OF ARGENTINA IS ---.

LANDLOCKED

YOUR ANSWER IS INCORRECT.
THE ANSWER IS: ATLANTIC.

THE SOUTHERN COAST OF ARGENTINA IS THE ANTARCTIC.

Fig. 1 Excerpt from dialogue between high-school student and SCHOLAR in Tutorial mode. (Student responses are enclosed in asterisks.)

THE WESTERN BORDERING COUNTRY OF ARGENTINA IS ---.

◆ I DON'T KNOW. ◆

THE ANSWER IS: CHILE.

THE BOUNDARY OF THE WESTERN BORDERING COUNTRIES
OF ARGENTINA IS THE ANDES.

THE NORTHERN BORDERING COUNTRIES OF
ARGENTINA ARE BOLIVIA AND PARAGUAY.

THE NORTHEASTERN BORDERING COUNTRIES ARE BRAZIL
AND URUGUAY.

THE CAPITAL OF ARGENTINA IS ---.

◆ BUENES AIRES ◆

I BELIEVE YOU MEANT TO TYPE "BUENOS AIRES".
VERY GOOD.

CORRECT OR INCORRECT?

BUENOS AIRES IS IN EASTERN ARGENTINA.

Fig. 1 (cont.) Excerpt from dialogue between high-school
student and SCHOLAR in Tutorial mode. (Student
responses are enclosed in asterisks.)

```

(RPA00
ARGENTINA
(((XN ARGENTINA)
  (DEF 2))
  NIL
  (SUPERC NIL COUNTRY)
  (SUPERP (I 6)
    SOUTH\AMERICA)
  (AREA (I 6)
    (APPROX NIL 1200000))
  (LOCATION NIL (IN NIL (SOUTH\AMERICA NIL SOUTHERN))
    (LATITUDE (I 6)
      (RANGE NIL -22 -55))
    (LONGITUDE (I 6)
      (RANGE NIL -57 -71))
    (BORDERING\COUNTRIES (I 2)
      (NORTHERN (I 1)
        (SEX BOLIVIA PARAGUAY))
      (NORTHEASTERN (I 1)
        (SEX BRAZIL URUGUAY)
        (BOUNDARY
          (I 2)
          URUGUAY\RIVER))
      (WESTERN NIL CHILE (BOUNDARY NIL
        ANDES))
      (SOUTHERN (I 3)
        CHILE))
    (COAST (I 1)
      (EASTERN NIL ATLANTIC)
      (SOUTHERN (I 1)
        ANTARCTIC)))
  (POPULATION (I 2)
    (APPROX (I 3)
      24000000)
    (ORIGIN (I 6)
      (PRINCIPAL NIL EUROPE)
      (COUNTRIES (I 2)
        (PRINCIPAL NIL ($L SPAIN ITALY))
        (RACE (I 6)
          WHITE
          (COMPOSITION (I 3)
            (WHITE NIL 95))))
      (LITERACY (I 4)
        90)
      (LANGUAGE NIL SPANISH)
      (RELIGION (I 2)
        (($L PRINCIPAL OFFICIAL)
          NIL CATHOLICISM)
        (OTHER (I 4)
          ($L JUDAISM PROTESTANTISM)

```

Fig. 2 Entry for Argentina in SCHOLAR's data base

Reentering the node Argentine, SCHOLAR searches for the next most important item, and selects location. The material embedded within Location is searched and SCHOLAR selects (IN NIL (SOUTH AMERICA NIL SOUTHERN)) to construct the question, "True or false? Argentina is in southern South America?" The student indicates true, which is the correct answer, and SCHOLAR affirms the correctness of the response. SCHOLAR next constructs a question about the eastern coast of Argentina. The student answers incorrectly, and SCHOLAR provides the correct answer and selects some related additional information within location, in this case, information about the southern coast, and presents it to the student. All of the presented material, as well as the material which was selected for the question, is tagged to reflect the first way in which it was used during the interaction. The answer to the question about the western bordering country was not known by the student, and, using the same strategy as before, SCHOLAR provides the correct answer and some additional related material.

SCHOLAR's strategy involves pushing down further into the topic until the student fails to answer a question correctly. SCHOLAR now pops up one level and selects the information about the capital of Argentina, to construct a question. Because the student answered the question correctly, SCHOLAR then proceeds to discuss aspects of "Buenos Aires."

Also in this mode, student questions of the type, "What is X?", "Where is X?", and "Tell me about X." are permitted (see an example in Fig. 3). About one-half of the total session (see experimental procedure) is spent reviewing from the top, reasking questions missed in the past and asking about material presented, as well as presenting new material. Figure 4 shows the reviewing of material covered earlier, shown in Fig. 1.

IS THE FOLLOWING TRUE OR FALSE?

CHILE IS IN SOUTHWESTERN SOUTH AMERICA.

◆WHERE IS CHILE◆

THAT WAS MY QUESTION.

CHILE IS IN SOUTHWESTERN SOUTH AMERICA.

Fig. 3. Example of a high-school student asking SCHOLAR a question in Tutorial mode. (Student responses are enclosed in asterisks.)

THE EASTERN COAST OF ARGENTINA IS ---.

ATLANTIC

THAT'S FINE.

THE SOUTHERN COAST OF ARGENTINA IS ---.

*ONE MINUTE ... *ANTARTIC*

THAT'S RIGHT.

THE WESTERN BORDERING COUNTRY OF ARGENTINA IS ---.

PERU

WRONG.

THE ANSWER IS: CHILE.

THE NORTHERN BORDERING COUNTRIES OF
ARGENTINA ARE BOLIVIA AND PARAGUAY.

THE NORTHEASTERN BORDERING COUNTRIES ARE BRAZIL
AND URUGUAY.

THE BOUNDARY IS THE URUGUAY RIVER.

IS IT TRUE THAT BUENOS AIRES IS IN EASTERN ARGENTINA?

T

VERY GOOD.

Fig. 4. Review of material shown in Fig. 1 by SCHOLAR in Tutorial mode. (Student responses are enclosed in asterisks.)

In Block-Test mode, the initial topic is again chosen from the prespecified agenda. SCHOLAR begins by selecting about eight facts from the most important unused material within the specified X level. For example, upon entering the node for Argentina, SCHOLAR selects some material, in order of importance, and presents it as shown in the sample dialogue (see Fig. 5). Immediately thereafter, SCHOLAR begins the questioning, about only those items presented. In the sample dialogue the student is asked about the superordinate of Argentina first, the location next, and so on, in the same order as that of the presentation. Once the answers to the questions are evaluated, SCHOLAR merely indicates the correctness of the response, never providing any additional information. The correct answer, however, is provided. Whenever the presented information contains a potential subtopic, it is added to the push-down list of topics to be selected. In the sample dialogue the presented paragraph contains mention of the Andes as the boundary of the western bordering country. The Andes, therefore, would be added to this list of subtopics to be covered upon the completion of the present line of questioning. Following that subtopic to completion, SCHOLAR, if no other topics were added within this topic, pops back up to the major topic, in this case Argentina, and proceeds. No student questions are permitted in this mode.

2.3 TESTING

A list of instructions was prepared for both modes, including a description of acceptable syntax for the student's questions to Tutor-SCHOLAR. Pre- and post-tests were prepared for each of the four countries. The pre-test was a randomly selected set of about half the questions from the post-test. It was used to assess the student's prior knowledge of each country, and served to equalize

ARGENTINA IS A COUNTRY.
IT IS IN SOUTHERN SOUTH AMERICA.
THE EASTERN COAST IS THE ATLANTIC.
THE SOUTHERN COAST IS THE ANTARCTIC.
THE WESTERN BORDERING COUNTRY IS CHILE.
THE BOUNDARY IS THE ANDES.
THE NORTHERN BORDERING COUNTRIES ARE BOLIVIA
AND PARAGUAY.
THE NORTHEASTERN BORDERING COUNTRIES ARE BRAZIL
AND URUGUAY.

ARGENTINA IS ---.

♦A COUNTRY♦

YOU ARE CORRECT.

IS IT TRUE THAT ARGENTINA IS IN SOUTHERN SOUTH AMERICA?

♦YES♦

THAT'S FINE.

USE ONE OF THE FOLLOWING:

ATLANTIC
PACIFIC
ARCTIC OCEAN
ANTARCTIC

TO ANSWER THE QUESTION:

WHAT IS THE EASTERN COAST OF ARGENTINA?

Fig. 5. Excerpt from dialogue between high-school student and SCHOLAR in Test mode. (Student responses are enclosed in asterisks.)

◆ATLANTIC◆

VERY GOOD.

WHAT IS THE SOUTHERN COAST OF ARGENTINA?

◆ANTARCTIC◆

THAT'S RIGHT.

THE WESTERN BORDERING COUNTRY OF ARGENTINA IS ---.

◆I DON'T KNOW◆

THE ANSWER IS: CHILE.

SELECT AN ALTERNATIVE FROM THE LIST:

ANDES
BRAZILIAN HIGHLANDS
SIERRAS
ALPS

TO COMPLETE THE SENTENCE:

THE BOUNDARY OF THE WESTERN BORDERING COUNTRIES
OF ARGENTINA IS ----.

◆ANDES◆

YOU ARE CORRECT.

Fig. 5 (cont.) Excerpt from dialogue between high-school student and SCHOLAR in Test mode. (Student responses are enclosed in asterisks.)

the subjects in each condition. Questions in both tests were carefully composed and sequentially arranged in a manner which did not give away information to the student. In addition to the geographical questions, the pre-test queried the student's exposure to South American geography, exposure to computers, and the subject's disposition on computers. Subjects were also questioned about their tendency to ask questions in school and on their preference for discussion vs. lecture classes, etc. These inquiries were made to ascertain if there might be an interaction between student preferences for controlling their own learning, and the amount of learning in Block-Test vs. Tutorial mode.

Each of the eight subjects were tested on all four countries, having learned about two in one mode and two in the other. The experiment was counterbalanced by country, day, and order. The questions on the post-test were separated into four countries (50 questions from each). Each country was divided into two separately administered sections because questions in the second part were likely to give away answers to part one. To the extent possible, questions were analogous from country to country. The pre-test consisted of 20 questions from each of the four countries (a subset of the 50) randomly mixed.

The sessions were conducted on successive weeks with students having two countries the first week and the other two the second week. (This was necessary because of the slow computer response on any day but Saturday.) The students were instructed as to the operation of the teletype terminals which they used and about the particular mode in which they were about to run. The sessions lasted about one hour for each country, with a five-minute break between successive runs on the same day. The students ran in both modes in one day. Students were asked to return on the Monday following the Saturday session each week for the post-test and

questionnaire. Subjects were local male and female high-school students who volunteered to participate in the exercise. The students indicated that they had had limited exposure in the past with South American geography, mostly in elementary school (5th to 7th grades), at a superficial level in the context of a world-geography class. Each student had at least some exposure to computers in their high-school math class where a mini-computer was available for student use. None of the students expressed any dislike or antagonism for computers.

2.4 RESULTS AND DISCUSSION

The difference scores between pre- and post-tests for each subject are shown in Table 1, broken down by presentation mode and by the order of the two modes on each day. To analyze the results of the experiment, we used a three way analysis of variance based on raw difference scores with mode, order, and subjects as the three factors. Since there was only one observation per cell, we took the mean square of the triple interaction as the estimate of error variance. Of the main factors, the effect of mode was significant ($F(1,7)=17.53$, $p<.01$), the effect of subjects was significant ($F(7,7)=14.45$, $p<.01$), and the effect of order was not significant ($F(1,7)=.38$). Of the two-way interactions, the interaction between mode and order was significant ($F(1,7)=10.58$, $p<.05$), and the other two interactions were not. (For subjects and mode, $F(7,7)=.73$ and for subjects and order, $F(7,7)=2.71$.) The significant interaction between mode and order reflects the fact that subjects remembered the second country they learned about on each day better than the first country. In this analysis, such a difference in retention shows up as an interaction.

TABLE 1

Raw Difference Scores between Pre-Test and Post-Test for each subject, broken down by Presentation Mode and the Order of the Two Modes on Each Day (Difference scores based on normalizing post-test scores by a factor of .4 are shown in parentheses)

Subjects	Block-Test Mode First		Tutorial Mode First		Average for Each Mode	
	Block-Test Mode	Tutorial Mode	Block-Test Mode	Tutorial Mode	Block-Test Mode	Tutorial Mode
S1	4 (1.6)	10 (4.0)	10 (4.0)	13 (5.2)	7.0 (2.8)	11.5 (4.6)
S2	2 (.8)	6 (2.4)	2 (.8)	1 (.4)	2.0 (.8)	3.5 (1.4)
S3	10 (.4)	18 (4.2)	14 (2.6)	16 (4.6)	12.0 (1.5)	17.0 (4.4)
S4	8 (.8)	16 (4.6)	8 (1.4)	5 (-.4)	8.0 (1.1)	10.5 (2.1)
S5	16 (6.4)	17 (6.8)	17 (6.8)	17 (6.8)	16.5 (6.6)	17.0 (6.8)
S6	12 (3.6)	19 (5.2)	10 (3.4)	16 (4.6)	11.0 (3.5)	17.5 (4.9)
S7	3 (-.6)	13 (4.0)	14 (5.6)	11 (3.8)	8.5 (2.5)	12.0 (3.9)
S8	6 (2.4)	10 (3.4)	11 (3.8)	13 (4.6)	8.5 (3.1)	11.5 (4.0)

The two effects we were interested in were the effect of mode, where Tutorial mode was clearly superior to Block-Test mode, and the lack of any interaction between mode and subjects. Taken together these two results indicate that the superiority of Tutorial mode was common to all the students and not just to those who prefer to control their own learning. Hence it is clear that of these two modes some aspects of Tutorial mode are of general benefit to student's learning of factual knowledge.

In general, students, when allowed to ask questions in Tutorial mode, did not ask SCHOLAR many questions. (In future work we will encourage them to do so more often.) On a questionnaire given with the final post-test, the students commented favorably about the Tutorial mode and particularly the procedure of going over material more than once. In contrast, they said Block-Test mode gave them too much information at once. Overall, students preferred the Tutorial mode over the Block-Test mode and indicated that they enjoyed reviewing questions that they missed. They also felt it was very helpful to get information related to the question they missed. Based on these comments, and the lack of questions by the students, the superiority of Tutorial mode probably was due to the reviewing in Tutorial mode and the excess of information presented at one time in Block-Test mode.

In the future we plan to use the method developed in this study to further explore what aspects of these tutoring strategies (and other variations) benefit students most. This is the fairest kind of comparison between teaching strategies, in and of themselves, because the other aspects of the teaching situation can be held constant in SCHOLAR. Our first attempt will be to compare

improved versions of both Block-Test mode and Tutorial mode. For Block-Test mode we will shorten the blocks and review questions within a block until the student answers correctly. In Tutorial mode we plan to cut down on the amount of additional information presented when an error is made and encourage students to ask questions more freely. Ultimately, we would like to make both Tutorial mode and Block-test mode as effective as possible so that students can choose, given their own preferences, which presentation strategy they want to use.

Another comparison we would like to make in a later phase of testing is an evaluation of the map display module now being added to SCHOLAR (see Section 4). To test the usefulness of maps, we would compare a version of SCHOLAR, which includes the map facility to an otherwise equivalent version without the facility. The pre- and post-tests could measure both map information and non-map information separately. It may turn out that students learn both kinds of information better with the map facility. That is to say, locating places visually on a map may help to tie in related, non-visual facts, so that they can be remembered more easily.

The fact that SCHOLAR can be used to test particular aspects of teaching methods makes it potentially a valuable tool for educational research. The possibility of trying out single modifications in teaching strategy to see their effects on students' learning rate is unique to SCHOLAR. Human teachers of course can make such modifications in their own teaching strategies, but there is no way to control all the other factors that might vary as they changed strategy. SCHOLAR, however, is in any specific version, a fixed system and so an unbiased comparison can be made using any number of subjects. After testing out single

modifications one at a time, it is possible to start combining those factors which show positive effects on students' learning, and to test them out in combination. In this way the accumulation of systematic knowledge about teaching methods can begin to occur.

SECTION 3

TEACHING PROCEDURAL KNOWLEDGE

3.1 INTRODUCTION

One of the goals of our work with the SCHOLAR system has been to study tutorial methods for teaching different kinds of knowledge. Because we had developed a second data base on the ARPA network for the Air Force (Grignetti & Warnock, 1973), the ARPA network was a natural context in which to study the teaching of procedural knowledge. Our basic approach is to study the strategies that good teachers use in tutoring procedural knowledge, and then to implement these strategies where possible in SCHOLAR.

The section describes our preliminary analyses of tutorial dialogues about the ARPA network; the decision to concentrate on NLS, which is a subsystem of the ARPA network, and to build a model NLS system within SCHOLAR; and our conclusions from the first few tutorial dialogues we collected on how to use the NLS system.

3.2 PRELIMINARY ANALYSIS OF DIALOGUES ON THE
ARPA NETWORK

In order to look at the sorts of problems which arise when attempting to convey procedural information, three tutorial sessions concerning the ARPA network were held between an experienced network user and a naive student. One session covered general information about the network and its usage, while the other two covered specific information about how to use FTP, the file transfer protocol used to transmit files from one node of the network to another.

In all three sessions, tutoring was done entirely via typewriter terminals; the student's terminal was linked to that of the teacher in such a way that whatever was typed on either terminal appeared simultaneously on both. In this way a SCHOLAR-like environment was imposed on the tutoring process.

In the first two sessions the terminals were used merely as a means of producing a student-teacher dialogue. The information typed was all in the form of comments describing some procedural system, rather than instructions to be executed by a system. In the third session the student actually attempted to use the file transfer system, receiving directions from the teacher (via the link) at each step along the way. This mode of instruction in which the student can actively participate has several clear advantages over the more passive situation in which he merely receives information:

- (1) he remembers things better for having done them.
- (2) he finds out what he doesn't know by being faced with the problem of actually doing things, rather than just giving or receiving descriptions of how to do them.
- (3) the student and the teacher need to interact less because the system being executed interacts with the student, giving considerable information in the form of prompting or explanatory messages.
- (4) unusual responses of the system can be dealt with and explained at the time of their occurrence and need not be described in advance.
- (5) unexpected responses generated by student error can be treated similarly.

For these reasons it seems clear that the teaching of procedural knowledge can proceed most effectively when the student is actually executing commands and observing the results of his actions during the course of his instruction. Therefore, it was decided that any attempt to teach procedural knowledge interactively with SCHOLAR should include a capability for trying out what is learned on a model of the system one is learning about. The idea is for SCHOLAR to be sitting on top of the model system available for teaching or answering questions. Eventually we would want SCHOLAR to be able to watch what happens between the student and the system he is exercising, just as a tutor does.

3.3 CHOICE OF THE TNLS SYSTEM

Having decided on the mode of instruction, thought was given to the particular body of procedural knowledge to be used for this study. Programming languages, which provide perhaps the most obvious examples of the use of procedures, were rejected as being too complex a subject area for an initial experiment. Although many languages have a fairly small and simple set of instructions, it is not the teaching of the meaning and effect of these commands which presents difficulties; rather the task is to convey how the commands may be combined to represent an algorithm suitable for the solution of a specified problem. A similarly difficult task is that of attempting to determine the intent of a set of instructions which do not produce the desired results, so that suggestions about suitable modifications can be made.

Consequently, the decision was made to study the command language of a system, rather than a programming language. The commands of a system are usually simple to learn and yet fairly

powerful, so relatively few commands may be needed in order to achieve a desired goal, and the intent of the user is more easily observed and determined.

After considering several systems (including TELNET and FTP, major subsystems for using the ARPA network), we chose to study TNLS, the Typewriter version of the NLS system developed by Doug Engelbart et al. at SRI-ARC. The TNLS is useful for text manipulation and editing. TNLS was selected for several reasons: it has a rich command structure; it provides sufficient depth of complexity so that the user often passes through a series of states in the attainment of his goal; interest in TNLS, and its display counterpart DNLS, is growing and the problem of teaching people how to become proficient in the use of this complex system is receiving increased attention.

3.3.1 Implementing a Model TNLS

Since TNLS is a very large system, encompassing many subsystems, it was necessary to choose a subset of the available commands and features in building a model system. Choice of the subset was based on experience gained in using TNLS to produce an actual proposal, and on a careful consideration of all features described in the User Guide; those which could be dropped with little or no loss of power, and which were thought to be seldom used, were omitted from the subset.

It was decided to write the TNLS subset in BBN-LISP so that it could be easily accessed from SCHOLAR. Students could then interact with this model system, generating results indistinguishable from those obtained from interacting with TNLS (provided only commands from the subset were used), and could also interact with SCHOLAR in order to ask questions about how to proceed.

Work on the model TNLS system should be completed by December, 1973. The preparation of a TNLS data base, which will be needed for SCHOLAR to answer questions about TNLS, is being started under a follow-on contract with the Air Force and should be completed by March or April, 1974. At that point SCHOLAR should be operative as a question-answering system about TNLS which could be used either in isolation or in conjunction with actual execution of TNLS commands using the model system.

Initially, SCHOLAR will not be able to "see" the interactions of the student with the model system, but will merely be available to answer well-formed questions (i.e., those without relative clauses, anaphoric reference, etc.) whenever they are asked, with no awareness of the context. At that point we will implement an event memory to be integrated with the semantic network. This will allow SCHOLAR to be aware of the past and current state of the user, thus enabling more sophisticated processing of student questions.

There are a number of advantages to working with a model TNLS system written in BBN-LISP, rather than with the actual TNLS system, which is written in a little used language called L10. The model, because it is a subset, will be much smaller; it will co-exist more easily with SCHOLAR; it will be written so that various kinds of information about the state of the program and hence the state of the user can be maintained and easily accessed. This last point is of particular importance for future developments in which an attempt will be made, using the event memory, to build up a history of the student so that some picture of his level of knowledge and perhaps his intent

can be formed; such information is of great importance in determining how to interpret questions and in deciding at what level to answer them.

3.3.2 Teaching TNLS

Shortly after these decisions were made, a teaching team from SRI-ARC came to BBN to hold two introductory classes in the use of TNLS. Both sessions, one for persons accustomed to the use of TENEX and similar systems and the other for persons with little such experience, were tape recorded and transcribed so that both TNLS data and the methods of teaching it could be studied. Numerous discussions were held with the three teachers involved about the various kinds of problems students encounter as they try to learn the system, and many of the actual questions and troubles of the students in these two classes were noted down for further analysis.

(When the instruction was completed, a list of recommendations concerning various aspects of TNLS was prepared and a consultation with Doug Engelbart and his staff was held at SRI-ARC to consider proposed revisions of TNLS syntax, and to discuss other features of a new version of TNLS to be released in the fall. Further cooperative efforts between members of this project and the ARC staff are planned.)

In order to look more deeply into the kinds of problems students encounter in learning procedural knowledge in general and the TNLS system in particular, five more tutorial sessions were held concerning TNLS, one a conversation recorded on tape, and the other four done over linked terminals as described above. In all cases the student was familiar with the use of systems, but unfamiliar with TNLS.

3.4 ASPECTS OF TEACHING PROCEDURAL KNOWLEDGE

The protocols from the various tutoring sessions were studied to determine what sorts of teaching techniques were used to tutor procedural knowledge. The following general approaches were noted.

3.4.1 General Knowledge and Specific Examples

An early protocol involving the use of the file transfer protocol had shown the teacher consistently using the approach of answering a student's question in general terms and then immediately following this general answer with one specific to the particular case at hand, or with an example. This is illustrated by the following exchange:

S: What's "filename?"

T: The name of the file you want to retrieve--in our case
<LOADSTAT>LDINF.SAV.

S: I see. What are the conventions covering filenames?

T: On TENEX a file name is of the form:

DEVICENAME:<DIRECTORYNAME>NAME.EXTENSION;VERSION

For example, DSK:<LOADSTAT>LDINF.SAV;1

names a file on the disk device, in directory <LOADSTAT>
with name LDINF, extension SAV and version 1.

A similar example from a TNLS protocol (with a different teacher) is given below, in which a definition of a general term is given in increasingly specific detail:

T: A branch consists of a specified statement and all other statements which have the same source; that is, all of the statements whose statement numbers begin with the same characters as those of the specified statement. Thus, branch 2 consists [in this case] of statements 2, 2a, 2b, and 2c.

The following description of the ARPANET and how to access it from BBN has the same properties:

T: The ARPANET is a set of computers scattered about the United States. They are all connected together so that a person at one site may use any of the computers at another site. The way that you connect to another site is to ask TENEX to call a subsystem called TELNET for you. Do that now by typing "TELNET."

These examples indicate that human tutors realize that procedures may be explained at different levels of generality and specificity. Therefore, the general rule and the specific example are presented conjointly. The general rule gives the student a model from which to generalize, and the specific example tells him exactly what to do or answers his question precisely.

Heretofore, SCHOLAR's data base has been restricted largely to general information. Examples are stored as instances of concepts (e.g., names of different computer centers or systems), but there is currently no way to store an example for a complex entity like a procedure or a branch in NLS (see the second example above). One way to discuss such an entity in both general and specific terms would be to store the general form, and then instantiate at output each of the parts making up the general form. This is what was done by the tutor in the last part of the first example. He gave the general form and then repeated the form, substituting an example of each part. An alternative might be to store a specific example of the entire general form under the entry where the general form is stored. It may be necessary to use both techniques.

A related aspect of this problem is suggested by the second example above. There, the tutor answers in terms of a specific

example which they had been discussing. To do this, examples will need to be stored in such a way that SCHOLAR can explain different concepts (such as branch, nesting, etc.) with respect to the example. This requires a flexibility in storage that SCHOLAR does not now have, but which is so essential to good teaching that we think it necessary to develop.

3.4.2 Similarities and Differences

Descriptions of new material may be given in terms of similarities to and differences from "old" material with which the student is already familiar. If similar concepts are involved, the old information will be helpful to him in acquiring the new; however, his old information may be a hindrance if the differences are not pointed out as well. For example, the following warning about TNLS commands was given to a student known to be familiar with TENEX:

T: "Note that TNLS commands have a different convention from TENEX commands. In TENEX you may type as many characters as you like and the system will echo the remainder. In TNLS you are allowed to type one or two characters only, for the most part. These characters are the first letters of the command words."

Students may themselves indicate their knowledge of related information in the posing of their questions, while indicating a desire to know how the new information is similar or different. The following student question illustrates this point:

S: "Is the cursor [in TNLS] positioned on a letter or between letters as in TECO?"

Given a new student, the teacher may begin by asking about his familiarity with related material. Such questioning may be done partly to determine the student's level of experience, but is also done so that the teacher may know in what terms he should proceed to make new definitions and descriptions so as to point out similarities and differences appropriately. The following quote is from the beginning of the first TNLS tutoring session:

T: "Before actually attempting to use the system, I would like to tell you something about the file structure in TNLS, which is different from that of most other file-handling systems. Are you familiar with TECO, for example? [Yes.] Then you know that TECO understands about lines and that the lines are ordered, but that is about all the structure there is. TNLS files are structured like an outline; that is, they look as follows":

Here (and elsewhere) the contrast is implicitly rather than explicitly stated, but the point is that TNLS files are structured while other files are not.

At present, similarities may be expressed in the data base only in the sense that items having the same SUPERC (super-concept) or SUPERP (super-part) may be said to be similar. Such relationships may be much too broad for this purpose, and a new special attribute specifying similarity should be introduced. Such an attribute must permit the specification (by embedding) of the ways in which the two items are similar or different. A more explicit indication of similarities and differences can always be derived by comparing the data base entry of both items for common and contrasting properties, as a subroutine in SCHOLAR now does.

3.4.3 Partial Answers (Hints)

With students who have had some experience, a teacher will sometimes answer a question only partially, or give some sort of hint rather than a precise response. This is done to force the student to discover the answer for himself, in the hope that he will then remember it better in the future. Hints are used when the teacher feels that the student should know the answer because of his previous experience with the problem. One example from the dialogues is shown below:

```
STUDENT (to System): Insert Statement after A:.1
SYSTEM:                .1?
TUTOR:                 You've forgotten about insert AFTER
STUDENT (to System):  Insert Statement after A:.Ø
```

Here, the tutor could have told the student what he did wrong, or told him what to do to correct his error, but instead the tutor reminded the student that insertions are made in TNLS not at the position specified, but after the position specified. This hint was enough for the student to figure out what to do.

Implementation in SCHOLAR would revolve around the proposed event memory, from which a record of the student's history could be built. The kind of answer he received to a question could be determined by his familiarity with the subject, based on the number of times he had embarked on similar procedures.

3.4.4 No Answer (Try It and See)

Since the student will have access to the model TNLS system while he is learning, it will sometimes be appropriate for SCHOLAR not to answer his question at all, but to indicate that he should

be asking the question of the model system instead, i.e., that he should "try it and see." The following instance of such a response showed up in the fourth TMLS tutorial when the student was already fairly knowledgeable about the system.

S: If I deleted a statement and then asked for it to be printed by its SID number, would it [the system] know what I meant?

T: Try it and see.

Although it may be difficult for SCHOLAR to determine when such a response is appropriate, it is probably the case that questions which begin with "if" and then specify the execution of some command could all be helpfully answered in this way. The actual result produced by the system will be provided more quickly, accurately, and memorably than any description of such a response which could be provided by SCHOLAR.

3.4.5 Answering with Yet Another Question

If the degree of sophistication of the student is not known, the tutor may resort to answering a question by posing a different question, one to help him to form a model of the student's knowledge so that he may respond at the appropriate level. Norman (1973) describes this process as follows:

"When we teach someone else knowledge, we are trying to build within that person a data base comparable to that of our own for the particular subject matter of interest. But in order to do this we must know what the other person knows and what he lacks. What is needed is some sort of interactive process in which we first question the other person to find out what is lacking, then teach, and then question again to find out how successful we have been."

If information about similarities and differences were provided within the data base, as described in a previous section, then SCHOLAR

might use this approach of responding with yet another question somewhat as indicated by the simulated protocol below:

- S: What is the syntax of the copy command?
T: Are you familiar with the syntax of the move command?
S: Yes.
T: The syntax of the copy command is similar to the syntax of the move command. The only difference is that you type c for copy instead of m for move.

In this way, a great deal of helpful information can be provided with very little text. Besides, overtly pointing out the fact that the commands are virtually identical in form is a more useful thing to do than presenting the complete syntax of the new command, and allowing the student to make the discovery for himself that it is the same as something with which he is already familiar.

If the student answers "no" to its question, SCHOLAR might persist and ask yet another question if there were another entity with high similarity in the data base. If the student has no useful previous knowledge, then of course a complete answer to his question must be provided.

Norman continues:

"In answering a question, it is important to be able to do more than simply combine information about the world with information that has been learned about the question. In order to derive the proper answer we must determine exactly why the question was asked, else we are likely to answer at the wrong level. This means that in addition to the knowledge of the subject being asked about, we must also have knowledge of the person who has asked the question."

This is a much more difficult approach to implement since the question of intent is involved. To know the intent of the question,

it may be necessary to know the intent of a sequence of commands which the student has been executing. This sequence, stored in the event memory, may be compared with some standard sequences for doing certain tasks and inferences drawn accordingly; however, there may be many ways of reaching the same goal (although the variations are far fewer with a command language than with a programming language), so the intent may not be easily discernible. The problem is further compounded by that of unintentional commands which are executed; there are many examples in the protocols of students inadvertently striking the wrong key and causing unexpected changes to occur. One approach to the problem is to ask the student to specify the intent of his question ("Why do you want to know?"), or the intent of his action ("Why did you do that?"), but the problem of comprehending his response will be sizeable.

3.5 FUTURE PLANS

Further tutoring is planned, using students with different levels of experience, including some unfamiliar even with the use of a terminal. Data gathered from the two teaching sessions given here by people from SRI showed that the kinds of questions and problems which arose in the experienced group were very different from those which arose in the inexperienced group; more study of both problem areas is needed.

In an attempt to simulate SCHOLAR's initial inability to see what the student is doing, some tutoring sessions over linked terminals will be tried, as follows. The student will embark on a specific task and will link to the tutor whenever he needs to ask a question. The tutor will have no knowledge of the student's actions, so the student will be forced to provide enough information in his question to obtain an appropriate answer. When his needs have been satisfied, the student will break the link and proceed once

more with his task. Such a simulation should be very helpful in determining the sorts of information SCHOLAR will need to possess in order to be useful in this situation.

As described above, a model TNLS system will be written in BBN-LISP and a TNLS data base will be formed. SCHOLAR can then be used as a question-answering system which a student can interrogate while actually exercising a model TNLS system.

A primer will be written during the next year, which will be used to introduce beginning students to the most basic aspects of the TNLS subset. This primer will be used by the tutors in the generation of protocols, and will no doubt be modified as experience with teaching TNLS is gained.

When the primer has been tested sufficiently and found to be a productive teaching aid, it will be implemented in SCHOLAR. The topics of the primer will be specified on SCHOLAR's agenda and eventually a student should be able to work through the primer, executing his commands in the TNLS subset while SCHOLAR "watches" to see what he is doing. SCHOLAR will present information as each new topic on the agenda is reached and will instruct the student to do some standard tasks. The student should be able to interrogate SCHOLAR at any point when he runs into difficulties using the subset, or has a general question about the TNLS system.

The problem of intent remains a large one, but is somewhat reduced in this environment since the intent is presumably to perform the specific task which has been assigned. This is a long-range project involving an improved English comprehension system, a complicated event memory, a more sophisticated semantic network capable of representing examples and containing new relationships, such as similarity. Much basic research into the problem of intent will be needed.

SECTION 4

TEACHING GEOGRAPHY WITH MAPS

4.1 INTRODUCTION

In the past few months we have been implementing a system to generate questions and evaluate answers concerning maps in SCHOLAR. When completed, this system together with the question answering module developed earlier under a different ONR contract (N00014-70-C-0264), will provide a mixed-initiative display system for SCHOLAR. At that point we will tie this system to Tutorial mode in SCHOLAR so that it can present map-related material as well as generate questions about maps. When the three systems are completed, SCHOLAR will be able to combine graphical and verbal information in teaching geography to the student. We then plan to use this system in further evaluative experiments (see Section 2) of the SCHOLAR system.

The three primary subdivisions of the display question generating system are the following: Topic and question generation, answer evaluation, and student error diagnosis. The first module has been largely completed; the other two are currently being designed and developed.

4.2 TOPIC AND QUESTION GENERATION

Topic selection for the new graphics package utilizes the weighted random strategy in mixed-initiative mode of SCHOLAR. Eventually it will also be called by the topic selection routines in Tutorial mode. Once a topic has been selected, the appropriate map is chosen for display. For instance, let us say that the

selected topic were Lima, the capital of Peru. The internal display figure representing a capital city is two concentric squares. The map generating heuristics have to determine that the appropriate map to display is the map of Peru centered on the screen with the symbol for a capital displayed where Lima is relative to Peru's outline. The capital city symbol representing Lima may then be blinked, or intensified independent of the rest of the display, to focus the student's attention on it. When one particular country is displayed the borders of surrounding countries are also displayed at lesser intensity to aid the student in placing the country in the appropriate context within South America.

The types of questions that SCHOLAR will ask are based on the questions tutors used in tutoring South American geography (Collins, Carbonell, & Warnock, 1973). There were four basic types of map related questions, each of which was phrased in a variety of ways. They were as follows: (1) Point to X (e.g., "Where is Cape Horn?") (2) Name and point to the Y's in X (e.g., "Why don't we try to name each of the countries in South America?") (3) What is the Y of this X (e.g., "What is the climate in this region?") and (4) general questions where the map may help the student (e.g., "What countries border on the Pacific?"). We have implemented the first two types, and the fourth will occur naturally because the map display will always contain the object under discussion. The third type is more difficult, but may be added later.

The two questions shown below were generated by SCHOLAR in the display mode and illustrate the first two types of questions.

- 1) Please point at the delta of the Orinoco.
- 2) Please name the principal cities of Chile that you know, and point to them on the screen in the same order that you name them.

The display map generated for the first question is a map of South America with several important rivers on it; the map for the second question is an outline of Chile with the cities blinking inside it. (Figure 6 shows the display for this question.)

As an example of the fourth type of question, SCHOLAR might ask:

- 3) Is it true that the Parana River is navigable?

At the same time, SCHOLAR displays the Parana within the appropriate context of South America. The map may then help the student to infer an affirmative answer from the fact that the Parana flows through several important cities, and is quite long.

4.3 ANSWER EVALUATION AND ERROR CORRECTION

The answer evaluating module is currently being developed and implemented. In evaluating the student's answer, which may be composed of both verbal and visual map input, certain semantic matching heuristics have been designed to simplify the otherwise monumental task. The student's verbal and map input responses are analyzed together for semantic sense and following that for factual correctness. In answering question two, mentioned earlier, the student may correctly identify one or two cities, but may also name one city in Chile while pointing at another; further he could name cities outside Chile, or geographical names which are not cities. This is where it is helpful to make some "sense" out of the student's answer before analyzing for correctness and diagnosing the errors.

KYBD → TENEX

THE POPULATION OF CHILE IS APPROX --- PEOPLE.

I WOULD GUESS AROUND 10000000

APPROXIMATELY TRUE, MORE ACCURATELY THE ANSWER IS 9,000,000 PEOPLE.



Fig. 6 A hard copy from the map display showing a sequence of two questions generated and the accompanying maps

KYBD→TENEX

PLEASE NAME THE PRINCIPAL CITIES OF CHILE THAT YOU
KNOW, AND POINT TO THEM ON THE SCREEN IN THE SAME
ORDER THAT YOU NAME THEM.



Fig. 6 (cont.) A hard copy from the map display showing
a sequence of two questions generated and the
accompanying maps

The third phase of the system, error diagnosis of the student's answer is still in the design stages. As previously mentioned, there are many reasons why a student's answer may be wrong. It is much more constructive to the student if he is told just what is wrong with his answer and given the right answer to compare to his own than if he is just told "wrong" and proceeding to the next topic. For instance, in the first example question where SCHOLAR requests the student to point at the delta of the Orinoco, let us say he pointed at the estuary of the Amazon. It would be useful for SCHOLAR to say the following:

You pointed at the mouth of the Amazon instead of the mouth of the Orinoco.
The mouth of the Amazon is an estuary, not a delta.
The difference between an estuary and a delta is:
A delta has many branches, but an estuary is a wide mouth where fresh water and salt water mix.
This is the delta of the Orinoco.
[Here SCHOLAR displays and blinks the delta to call attention to it.]

Error analysis should discover the dual error; the student pointed at the wrong river, and confused a delta with an estuary. We hope to eventually have an error analysis system that can generate the above output, although at present we are still in the flow chart state pending completion of the answer evaluating module.

SECTION 5
REFERENCES

- Carbonell, J.R. AI in CAI: An artificial-intelligence approach to computer-assisted instruction. IEEE Trans. on Man-Machine Systems, MMS-11, 190-202, 1970.
- Collins, A.M., Carbonell, J.R., and Warnock, E.H. Analysis and Synthesis of Tutorial Dialogues. Technical Report for Contract No. N00014-71-C-0228, September, 1973.
- Grignetti, M.C. and Warnock, E.H. A Mixed-initiative information system for computer-aided training and decision making. Final Report for Contract No. F19-628-72C-0163. September, 1973.
- Norman, D.A. Memory, Knowledge, and the answering of questions. In R.L. Solso, Contemporary issues in cognitive psychology: The Loyola Symposium. New York: Halsted Press, 1973.
- Warnock, E.H. and Collins, A.M. Research on semantic nets. BBN Report No. 2555, May 1973.