Instructional Science 6 (1977) 135–149

REAL-TIME SIMULATION IN COMPUTER-ASSISTED INSTRUCTION*

STANLEY TROLLIP and ANDREW ORTONY

University of Illinois at Urbana-Champaign, Illinois

ABSTRACT

Large scale computer-assisted instruction systems generally impose severe constraints upon the demands that individual users may make. Nevertheless, it is possible to overcome these problems and a program is described which teaches students how to fly a specific maneuver through real-time simulation of the flight of a student-controlled "airplane". This is achieved in spite of the fact that the student is "flying" his "plane" through the use of a manually controlled analog input device. Both computational and educational implications are discussed.

Introduction

Computer-assisted instruction is coming of age. Research and development are widespread and the range of applications is constantly expanding. Yet, in many ways a certain conservatism dominates the field – a conservatism resulting largely from a lack of sufficiently flexible systems. This paper describes an example of the way in which one might break away from the rather routine presentation of educational material which is still so prevalent. It describes a teaching program which utilizes the full potential of the computer to produce for the student a different kind of learning experience – an experience which allows him to acquire not factual material, but a skill. The program is an example of a computer being used for encouraging "knowledge how" rather than "knowledge that".

In normal instruction much of what is learned is learned as a series of rules or principles, and most computer-assisted instruction (CAI) has capitalized on the fact that computers are well suited to the branched presentation of material of this kind. Although CAI programs do exist that teach typing, electronic circuit design and debugging, and music composition, the

* Contractual support for this work was provided in part by the Life Sciences Program Air Force Office of Scientific Research, Contract number F44620–70–C–0105. Dr. Charles Hutchinson was the scientific monitor of the contract.

acquisition of skills has generally been left to interaction with the real world. The program we describe simulates the relevant aspect of the real world and allows the student to learn something which is inherently resistant to being explicated in an explicit rule-like form. Although the particular program we describe makes use of an analog input device, the important general issues are, we believe, neither affected by nor dependent on such special-purpose devices.

The particular problem to be discussed concerns teaching pilots effective strategies for flying holding patterns in varying wind conditions. Traditionally, this is a difficult and expensive problem: difficult, because wind necessitates changes in the shape of the pattern to compensate for drift; expensive, because suitable experience has hitherto only been possible through the use of either a simulator or an airplane. In our approach, an external analog input device, a hand-controller or "joystick," is interfaced into a large-scale general-purpose CAI system, in this case the PLATO system (Alpert and Bitzer, 1970) at the University of Illinois. The role of the joystick is to provide the student with controls similar to those he has in an airplane.

The Problem

A holding pattern is a maneuver designed to ensure that the ground controllers know exactly where each plane is in a high density traffic situation. Orders to hold are usually given when there are a number of planes waiting to land at the same airport, or when a plane has to wait for clearance to proceed on its course. It is a means of stacking many planes in a small area while ensuring maximum safety.

In a no-wind condition the holding pattern is racetrack in shape with the end of one of the straight legs defined by some radio fix. The pattern is flown with the radio fix marking the end of the inbound leg. There follows a 180° standard rate turn (3° per second), to the right unless otherwise specified. The size of the pattern is defined by making the inbound leg one minute's flying time, and the direction by specifying the course of the inbound leg. To ensure adequate safety, each aircraft is assigned protected airspace which is five miles wide on the holding side of the inbound leg, two miles wide on the non-holding side, and ten miles long. In addition, the altitude at which the pattern is to be flown is specified, with 1000 feet vertical separation between aircraft.

The task of flying holding patterns would be relatively simple if there were never any wind. In practice, however, this is rarely the case, and consequently, pilots need to be taught procedures to deal with windy conditions. As has already been mentioned, one of the determining features of the holding pattern is that the inbound leg be one minute long, and it is this feature that remains constant over all situations. If the inbound leg, whose direction is specified, is to remain one minute long irrespective of the wind, then the shape of the pattern will have to change in order to accomplish this.

For example, if a plane is flying at an airspeed of 80 miles per hour, and on the inbound leg there is a headwind of 20 mph, then the ground distance covered in the one minute on the inbound leg will be less than if there were no wind. In fact the length of the inbound leg would be exactly 1 mile (since the groundspeed is 60 mph), as opposed to 4/3 mile in the windless situation.

Continuing the example, since there was a 20 mph headwind on the inbound leg, then on the outbound leg there would be a 20 mph tailwind resulting in a groundspeed of 100 mph. Because the length of the outbound leg is approximately equal to the inbound leg, it will be about 1 mile long. At 100 mph the time on the outbound leg would be $(1 \times 60)/100$ minutes, or 36 seconds. Thus one procedure that is available to deal with different headwind components is to estimate the time it would take on the outbound leg to give an inbound leg of one minute. Initially the time on the outbound leg is estimated from knowledge of existing wind conditions and is improved each time around the pattern.

In crosswind situations the difference in groundspeeds occurs not on the inbound or outbound legs, but rather in the turns. Thus when turning into the wind groundspeed drops, and since the rate of turn remains constant at 3° per second, the radius of the turn decreases. When turning downwind, groundspeed increases resulting in a turn of a larger radius. In addition, in a crosswind situation it is necessary to head the plane into the wind on both the inbound and outbound legs to maintain the desired course. This is called "crabbing" into the wind and is similar to tacking in sailing. Figure 1 illustrates the pattern shape if there is a strong wind from the left on the inbound leg.

The important question for the pilot is how to determine the outbound heading in order to join top and bottom turns which are of different sizes. A good heuristic is to hold twice the inbound crab angle on the outbound leg. Thus if a 10° crab held the desired course on the inbound leg, then holding a 20° crab on the outbound should join the two turns. The rationale behind this heuristic is that one cannot crab while turning; therefore, by holding twice the crab angle on the outbound leg, one compensates for the lack of crabbing in the turns. In practice this works quite well.

As can be seen from the foregoing discussion, the techniques for compensating for different wind conditions are not excessively complicated. However, it is the experience of many flight instructors that performance on this maneuver is poor. Perhaps a reason is that traditionally pilots are taught

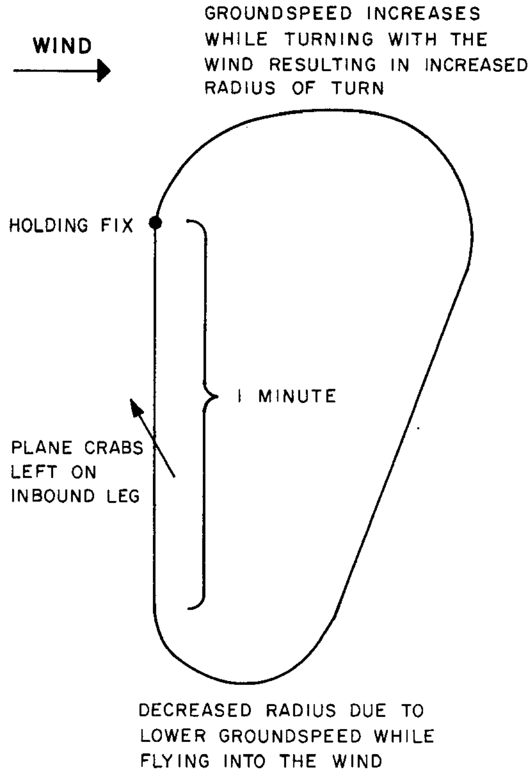

Fig. 1.

that the holding pattern is racetrack in shape with minor variations for different wind situations. Our approach is to emphasize that in most situations the holding pattern is *not* shaped like a racetrack, but is distorted according to the wind direction and strength. That is, emphasis is placed on how the shape of the pattern changes to ensure that the inbound leg remains one minute long. By learning the task in this way, the pilot should not only have a better understanding of the relationship of the plane to the ground in terms of speed and direction, but also be able to visualize better where in the pattern he should be and in which direction he should be flying. This latter point has interesting educational implications which will be touched upon later.

The Program

To accomplish its purpose, the holding pattern program provides the following capabilities. First, it has an external hand-controller and throttle linked into the PLATO terminal allowing the user to "fly" a simulated airplane around the screen. This capability requires that the hand-controller be able to signal to the computer left and right turns, and climbs and descents. Second, the program displays a scaled picture of the plane's progress relative to the ground. That is, the path drawn on the screen accurately represents the plane's groundtrack. Third, the basic instruments necessary for flying holding patterns are simulated so that the student can fly a precise pattern from them. These instruments include an altitude indicator, a heading indicator, an airspeed indicator, a turn indicator (which allows the student to estimate the rate in degrees per second at which the plane is turning), and some appropriate radio-navigation receiver (see example in Fig. 2).

Thus, typically, when a student logs onto the system, he is first interrogated by the computer as to his prior experience, such as ratings,

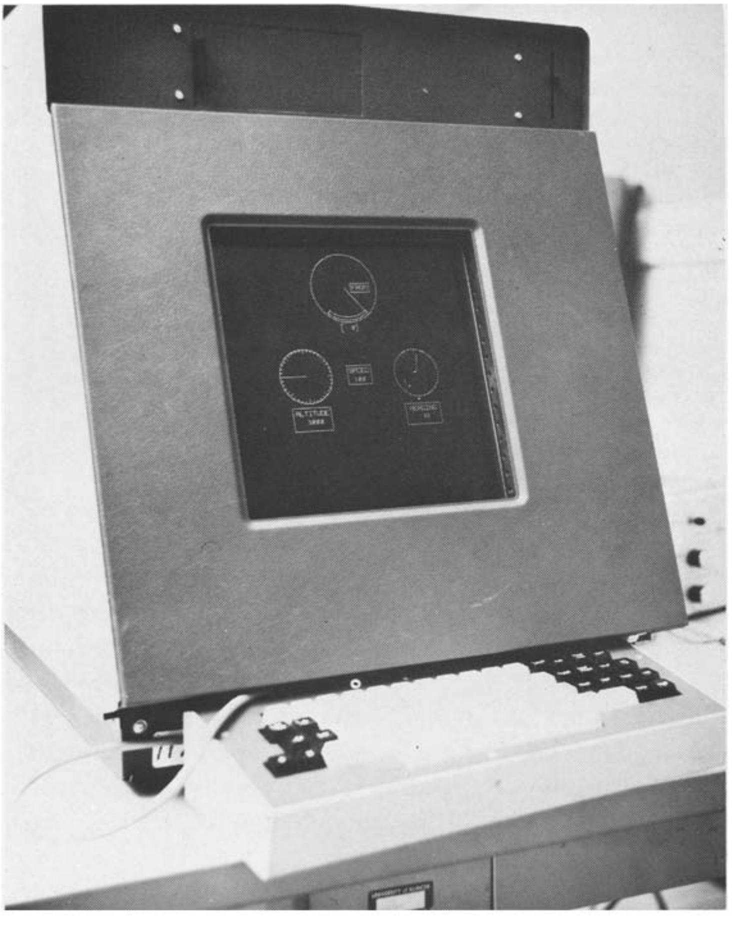

Fig. 2. The PLATO terminal and keyboard with the flight instruments displayed.

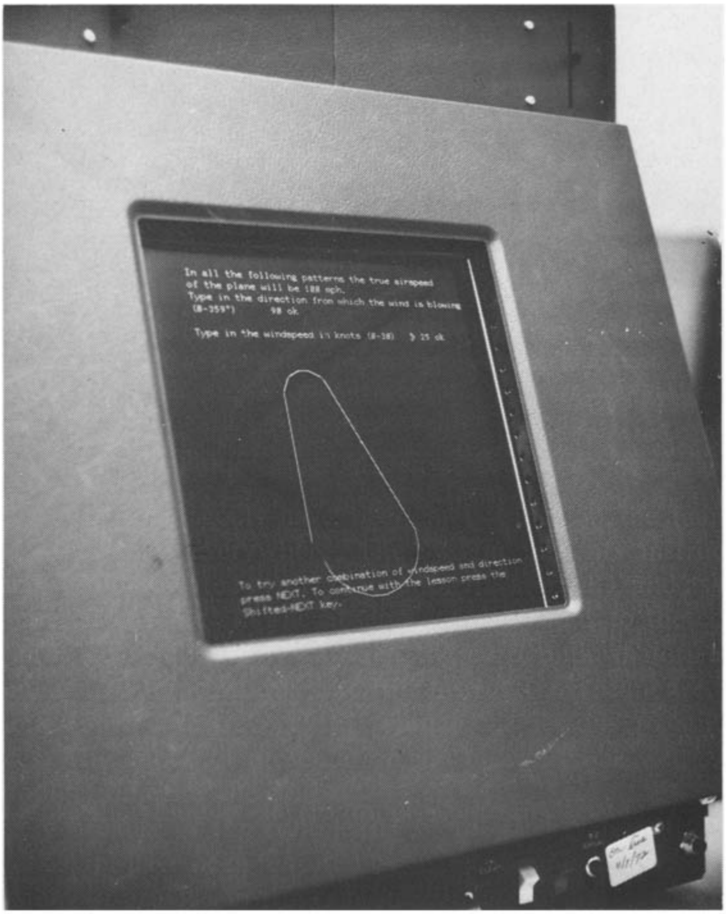

Fig. 3. The student can have the computer generate the ideal path for any wind condition. Here he has chosen a right 25 knot crosswind.

recency of experience, and total hours flown. This information is used later in the program to set the initial difficulty of task by using appropriate wind conditions, or by "freezing" the altitude or rate-of-turn apsects of the flight. The student is then given the chance to select various wind directions and speeds and is allowed to examine the shape of the "ideal" path to which these selections give rise (see Figs. 3 and 4). Finally, the appropriate instruments are displayed and the joystick activated, and he has to fly holding patterns until criterion has been reached. On-line feedback is available on his current, as well as his overall, performance.

In the traditional CAI lesson, the criteria on which student performance is assessed are usually quite simple and readily available, generally taking the form of answers to multiple-choice questions. In such cases the question

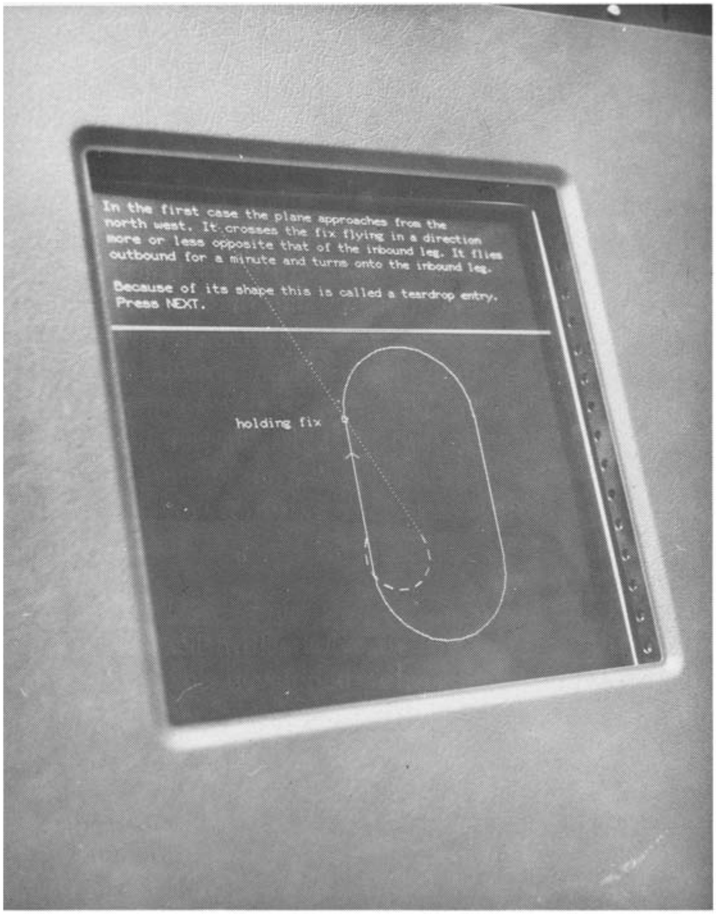

Fig. 4. A portion of the instructional sequence on entry procedures.

plays the role of an error sensor (Crowder, 1960). Since there are usually a very small number of possible choices to any question, it is relatively easy to decide what to do with the student on the basis of his answer. In some of the more advanced programs, branching to remedial routines is based on the answers to more than one question, and remediation may be invoked only if the student has less than a certain percentage of all responses correct. Smallwood (1962) went a stage further and not only determined the path of each student through the program on a basis of immediate performance, but also used the performances of all who had already taken the program in the branching strategy.

The holding pattern program requires a different approach to student evaluation since the task is continuous and not discrete, both in terms of time and the range of responses. The range of possible errors is greater in the sense that the identical performances may result from one or several different performance deficiencies. For example, if the inbound leg were flown consistently to the left of course, possible reasons for this could be incorrect compensation for crosswinds or perhaps incorrect use of the radio navigation equipment.

Since there are several possible causes for each incorrect position during the task, a spontaneous diagnosis of mistakes is unlikely to be very accurate or helpful. Consequently, a meaningful evaluation of overall performance can only occur after an analysis of data collected over a sustained period of time, such as two or three full patterns. This we call the after-flight feedback. However, during the "flight" a second level of feedback, the in-flight feedback, is necessary. This takes the form of providing information as to whether the plane is to the left or right of course, whether turns were initiated on time, or whether time on the outbound leg was chosen correctly to result in a one minute inbound leg. The in-flight feedback can be likened to that given to a student by the instructor as the task is performed, while the after-flight feedback is equivalent to the flight debriefing.

In its present form, the in-flight feedback is optional, allowing the pattern to be flown both in an instructional setting with feedback, or in a testing mode, without. In the testing mode, the program suppresses the information displayed on the screen, but nevertheless stores it for later remedial use. As well as providing the student with appropriate training, it is important that the program, like an instructor, knows when it can be reliably concluded that the student can fly holding patterns safely. This is achieved in the following manner. Before the student starts flying the first pattern, the program calculates how long the top turn should take, so that the roll-out is on the appropriate heading to compensate for drift. Once the plane crosses the holding fix, a clock is started. When the calculated time has elapsed, the plane's position is recorded, and the vector between this position and the calculated position is computed. Students who are within the permitted area are deemed to have demonstrated competence, while all others are required to continue until the criterion is satisfied. This critical region was determined by allowing many pilots of varying ratings and capabilities to fly the program's plane, recording their deviation from the desired point. From this data a value was selected that seemed to be a realistic index of competent performance. The reliability of this index can only be determined over time by assessing how well pilots fly holding patterns in a real airplane, having successfully reached criterion on our program. Should the level of performance of these pilots be too low, the criterion can be made more demanding. This empirical validation of the program's criterion is extremely important, if the program is to be considered for extensive real-world use. Initial results indicate that this is an effective measure of performance.

A further way that this measure of performance can be used in the program is as a basis for an adaptive teaching strategy. A series of nested critical regions could be empirically established that would indicate not only final and acceptable overall competence, but also perhaps competence at various subtasks. Thus, the program could adapt the difficulty of the task to the demonstrated proficiency of the pilot by varying one or more of the contributing variables. A novice or student pilot may be required to fly a holding pattern with the program automatically holding constant such factors as altitude and rate of turn. An experienced pilot may have to fly his pattern in difficult wind conditions, with the program providing random turbulence affecting rate of climb and rate of turn. In this way, the pilot is always flying at a level within possible reach of his current capabilities. This gives him the opportunity of mastering each aspect of the task without being distracted by other problems. As his proficiency increases, the difficulty of the pattern to be flown adapts accordingly, until finally the overall criterion is satisfied.

To provide the necessary error diagnosis, the program records the mean deviation of the flown course from the "ideal" course on the inbound and outbound legs, but not on the curved sections of the pattern. Also, the mean headings on both the straight legs are recorded to check whether effective wind correction angles were used.

The assessment of performance on a continuous task, like the one under consideration, is very difficult to do automatically. Even in the real situation of flying an airplane in a prescribed maneuver, there is considerable difference of opinion even among flight instructors as to how one objectively assesses performance. One advantage that the program has over the real situation is that it knows where the plane is relative to the ground at any moment, while in reality a pilot does not. The program can make use of this unique capability both as the determining parameter in an adaptive model and in the evaluation of student performance. Thus, in some respects the program has more ready control over the student than a flight instructor would have.

The Implementation

It should by now be evident that a program of the kind we have described would present no problem if it were implemented on a more or less dedicated machine, for, although many of the details of its implementation are complicated, the program makes no special demands in terms of novel hardware or software. However, what we take to be the important issue is not just the program itself, but the relation between such a demanding real-time simulation and a heavily-used general-purpose CAI system.

Consequently, we propose to discuss this relationship with special reference to the limitations imposed on the program by the generality of the system. Although this discussion will have to be in terms of the PLATO system on which the program is implemented, we would hope that our conclusions are of more general significance.

The PLATO system is a general purpose CAI system although it differs from other existing systems in a number of ways. PLATO was conceived, designed, and built as a teaching device. Experts from the education, economic, and engineering fields collaborated in its design and have monitored its development through four generations of systems. It is both large scale and user oriented; it was designed to operate 1000 terminals all connected to the same central computing facility with a maximum response time of no more than half a second; and it was designed so that users can either program their own lessons with relative ease,* or change lessons provided by different curriculum groups. Program editing features are very powerful and many of the commands in the high-level programming language, TUTOR, have been determined by user demand.

Because of the complexity of the interaction in the holding pattern program in terms of input and output, both the hardware and the software have to be regarded as potential inhibitors of the smooth operation of the program.

---

* Below is an example of TUTOR code. (Adapted from Ghesquiere, Davis, Thompson, 1974.) Explanatory comments follow $$ sign.

```
unit    gorge     $$ Defines name of small portion of program.
at      1410      $$ Specifies where next screen writing is to appear.
write   Where is Louis S. B. Leaky's anthropological dig?
                  $$ This appears on screen.
arrow   1610      $$ Readies PLATO to expect student response
                  $$ which will appear at location 1610.
specs   bumpshift, okspell $$ Permits upper and/or lower case
                  $$ letters to be input. Also permits minor spelling
                  $$ errors to be accepted as correct.
answer  < the, it, is, in, at, kenya > olduvai (gorge, canyon)
                  $$ Words between < > will be ignored,
                  $$ olduvai is mandatory, and
                  $$ either 'gorge' or 'canyon' is acceptable.
write   Homo habilis was discovered there.
                  $$ This appears on screen if the correct
                  $$ answer is given.
wrong   < the, it, is, in, at, tanzania > gombe stream research center.
                  $$ This anticipates a particular wrong answer.
write   That's the site of Jane Goodall's work with chimpanzees.
                  $$ Appropriate response for particular wrong answer.
no                $$ Catchall for all other incorrect responses.
write   Wrong   Try again. $$ Comment for unanticipated response.
```

Although at first sight the data transmission requirements of the joystick are very high, these can be brought within reasonable bounds by transmitting only when it is moved. In this way all the necessary information can be transmitted in spite of the narrow bandwidth of the telephone line connecting the terminal to the computer. Apart from this, the system hardware imposes no limitations primarily because the display does not require constant refreshing. On the other hand, the system's software imposes more serious constraints.

One of the fundamental problems that has to be solved in any time-sharing environment is the distribution of processing time to meet the needs of every user. In a CAI system it is a reasonable assumption that the vast majority of users can be satisfied by a modest allocation of time, because first, most users are thinking most of the time, and second, because traditional CAI applications require relatively little computation. Just how much use each terminal may make of the central processor depends primarily on the number of terminals in operation and the power of the central processor. For example, on the PLATO system, which is based on a CDC Cyber 70, a program like ours which averages 4000 instructions/sec processor usage will not be taking an unfair share with terminals expecting a maximum response time of less than half a second.

As the number of terminals being simultaneously serviced increases, sacrifices will have to be made either in terms of the average amount of processing time each can use, or in terms of the average response time. And it is safe to assume that most CAI systems would opt for the first as the lesser of two evils. Consequently, as demand increases and these operating conditions become more stringent, it may become necessary to use real-time simulation programs in off-peak periods. In terms of the figures given above, however, which are typical for PLATO, the holding pattern program can still operate during peak periods without detriment to the system or the program.

When it ceases to be the case that a real-time simulation can operate within the constraints of a CAI environment, an attractive solution would be to augment the system with an "intelligent" terminal (Stone, et al., 1974). This is a terminal that is attached to a mini-computer, which in turn is connected into the large system. In this way the mini-computer essentially relieves the central computer of excessive time and space demands, and increases effective input/output bandwidth. The central computer is retained for large memory tasks, storage, user monitoring, and general performance evaluation.

Education Implications

The rationale for the development of the holding pattern program is rooted in three different areas. First of all, the project was undertaken because it was felt that it would be cost effective, competing successfully with alternate forms of instruction. Second, the "learning by doing" philosophy as expounded, for example, by Bruner (1973), and Papert (see Papert and Solomon, 1972; Papert, 1972a; Papert 1972b) could be incorporated into such a training program. Third, initial efforts in the uses of imagery techniques in training as reported by Prather (1973) could be extended and their effectiveness investigated.

In either a plane or a simulator, some considerable portion of the time is used, not by practicing those procedures that need practice, but rather by completing mundane and relatively unimportant procedures which do not require extensive rehearsal. For greater efficiency both in terms of time and money, CAI programs, like the holding pattern program, have a great deal of potential. Not only is the operation of such programs likely to be substantially less costly than even a simulator, but there is also a reduced need for an instructor, since there are no safety hazards in using a computer terminal and the program can offer feedback which is frequently comparable to that given by an instructor. (The cost per hour for a single-engine trainer is about \$ 20; for a simulator about \$ 8; and for PLATO about \$ 1). In a group setting, one instructor could monitor on a master terminal the on-going performance of several students at once.

However, the fact that the student no longer has to do all the things involved in flying the airplane, such as handling communications, navigation and area clearance, has greater implications than merely saving time and money, for it allows him to provide full attention to the task at hand, rather than having his attention divided by these other tasks. After all, the purpose of the program is not to teach a person how to fly a plane, but to teach him to understand how to fly holding patterns or related maneuvers. When these procedures have been mastered, then it is time to put them into practice. This isolation of the pertinent task is one of the powerful capabilities of CAI systems in general and has been used with considerable success in diverse areas (see Smith, 1970).

The holding pattern program allows the student to familiarize himself with the dynamics of the holding pattern situation; he learns what happens when the wind changes and how he should compensate for it, and practices what he has learnt in a non-threatening environment with the computer giving appropriate feedback. Thus, the student knows what it is all about when he gets into the plane and is asked to perform the maneuver. He does not have to concentrate on keeping the plane flying and at the same time try to work out how the pattern should be flown. More attention can now be given to the task of keeping the plane flying safely.

This use of CAI represents a dramatic change from the traditional application. Normally CAI lessons have been taught by merely providing the necessary information, sometimes in a manner that could not be done either by a teacher or a book. Anything learned would be as a result of the student digesting the information and remembering it. On the other hand, the holding pattern program teaches by allowing the student to experience and do what is to be learned. Looked at like this, the holding-pattern program actually simulates a simulator, whereas a traditional program would simulate a book or perhaps a teacher.

The third area which provided the background for the program was the use of imagery techniques to improve performance. Prather (1973) used a type of imagery instruction or mental practice as an adjunct to normal training to teach students how to land. In his study Prather played voice recordings of where the plane was in the landing pattern and what should be done to the controls moment-to-moment. Over the training sessions, the amount of detail in the tapes diminished. So in the first session the instructions were comprehensive, including desired airspeeds, throttle settings, altitudes, etc. However, in the final session the recordings would merely be, "You are now on downwind" and "You are now on final". The students sat in a cockpit mock-up and had a throttle and a control stick to manipulate but no working instruments to read. What was being taught was a mental representation of the landing procedure, so that the students knew what to do by referencing this representation. This is similar to the subjective experience of "keeping ahead of the plane". In the evaluation of Prather's techniques, the experimental group landed better than their control counterparts, who learned conventionally.

A similar technique was used by Feurzeig (1971) in two computer-based applications. The first, which provided much of the inspiration for this program, was also a holding-pattern program. The second was designed to teach the skills involved in estimating the relative courses of ships moving at different rates and in various configurations. In the evaluation of their programs, Feurzeig and his associates found that the instructional programs were effective.

A more abstract approach, which was based both on the notion that the student should have an appropriate mental picture of the task, and that this was best obtained by making the student an active participant in the learning process, was suggested by Goldstein (1972). He made use of the TURTLE in the LOGO project (see Abelson, et al., 1973) to let prospective pilots program different situations related to flying. "The goal is to provide a better environment for a pilot to build mental models of his plane's performance under different flight conditions and graphically explore situations he could never safely be exposed to in the air" (Goldstein, 1972). Making the student actually program the TURTLE to "fly" in different wind situations increases the probability that he will acquire a complete understanding of the concepts involved, while the output of the programs allow him to examine a picture of what happens under the different conditions. In this way the necessary abilities for adequate "thinking ahead" of the plane are provided.

The examples described above indicate that the use of this type of imagery training could lead to improved performance in two ways. First, the pilot will be provided with a readily accessible and vivid reference as to his performance if he keeps in his mind's eye a picture of where his plane is in the pattern and how it is progressing relative to the ground. Thus, if his attention is temporarily distracted by having to compensate for some error, it will be easier for him to regain orientation. The second improvement of performance should be apparent in the pilot's preplanning abilities. If the task to be performed is mentally visualized, then perhaps, this will of itself cause the pilot to "think ahead" of the plane. This latter ability of "thinking ahead" of the plane is perhaps the most important ingredient for safe flying. By being prepared for various contingencies, the shock of surprise of an unexpected occurrence is minimized, and the appropriate corrective actions are implemented.

The program is currently undergoing extensive evaluation. Initial results indicate that pilots in training react favorably to it. They have little difficulty in learning to fly PLATO's plane, and it seems that teaching the pattern shape as being dynamic rather than static offers them new insights into ground-referenced maneuvers. A final quantitative assessment of the program's value in training for holding patterns and its value for transferring this learning to different tasks will have to await completion of the evaluation.

Conclusion

The holding pattern program, therefore, has important educational implications, for if it is to be successful in training for adequate performance, then .it is most certainly going to be one of the cheapest media available for that training. Also with the proliferation of CAI systems, the accessibility of such training will become easier. Secondly, because CAI offers unrivalled capability for successively incrementing such things as attention requirements and approximations to reality, it allows thorough evaluation of the various techniques mentioned, and an assessment of what techniques work in different situations. Hopefully this work suggests that more radical uses of CAI offer great potential advantages. Through judicious use of real-time simulation and external input devices, CAI can be used for instruction in areas which transcend the acquisition of rules and concepts so common in current work.

References

Abelson, H.; Goodman, H.; Rudolph, L. (1974). *Logo Manual.* Cambridge, Mass.: Massachusetts Institute of Technology, Artificial Intelligence Laboratory, LOGO Memo, No. 7.

Alpert, D., and Bitzer, D. L. (1970). "Advances in computer-based education," *Science,* 167, 1582–1590.

Bruner, J. (1973). *Beyond the Information Given.* New York: Norton.

Crowder, N. A. (1960). "Automatic tutoring by intrinsic programming," in Lumsdaine, A. A., and Glaser, R., Eds., *Teaching Machines and Programmed Learning: A Science Book.* Washington, D. C.: National Education Association, Department of Audio-Visual Instruction.

Feurzeig, W. (1971). *Automated Instructional Monitors for Complex Operational Tasks.* Cambridge, Mass.: Bolt, Beranek and Newman, Report No. 2272, October.

Ghesquiere, J.; Davis, C.; Thompson, C. (1974). *Introduction to TUTOR.* Urbana, Illinois: University of Illinois at Urbana-Champaign, Computer Based Education Research Lab.

Goldstein, I., and Goldstein, G. (1972). *Flying in Your Head.* Cambridge: Massachusetts Institute of Technology, Artificial Intelligence Laboratory, LOGO working paper No. 3.

Papert, S. A. (1972a). "Teaching children thinking," *Programmed Learning and Educational Technology* 9, 5.

Papert, S. A. (1972b). "Teaching children to be mathematicians versus teaching about mathematics," *International Journal of Mathematics Education and Science Technology* 3, 249–262.

Papert, S. A. and Solomon, C. (1972). "Twenty things to do with a computer," *Educational Technology* 12, 4.

Prather, D. (1973). "Prompted mental practice as a flight simulator," *Journal of Applied Psychology* 57, 3, 353–355.

Smallwood, R. D. (1962). *A Decision Structure for Teaching Machines.* Cambridge, Mass.: Massachusetts Institute of Technology, Research Monograph.

Smith, S. G. (1970). "The use of computers in the teaching of organic chemistry," *Journal of Chemical Education* 47, 9.

Stone, M.; Bloemer, R.; Feretich, R.; Johnson, R. (1974). *An Intelligent Graphics Terminal with Multi-host System Capability.* Washington, D. C.: Procedures of COMPCON, Fall.