

DOCUMENT RESUME

ED 169 916

IR 007 202

AUTHOR Knerr, Bruce W.; Nawrocki, Leon H.
 TITLE Development and Evaluation of an Adaptive Computerized Training System (ACTS).--R&D Report 78-1.
 INSTITUTION Army Research Inst. for the Behavioral and Social Sciences, Arlington, Va.
 PUB DATE Sep 78
 CONTRACT DAHC-19-74-C-0027; DAHC-19-75-C-0013; DAHC-19-76-C-0019
 NOTE 44p.

EDRS PRICE MF01/PC02 Plus Postage.
 DESCRIPTORS *Artificial Intelligence; *Computer Assisted Instruction; Decision Making Skills; Individualized Instruction; *Job Training; *Military Training; Simulation; *Training Techniques

ABSTRACT

This report describes the development of a computer based system designed to train electronic troubleshooting procedures. The ACTS uses artificial intelligence techniques to develop models of student and expert troubleshooting behavior as they solve a series of troubleshooting problems on the system. Comparisons of the student and expert models provide the basis for the instructional feedback given to the student. The ACTS also simulates the circuit on which the student is being trained and the process of taking test measurements, so that no actual equipment is required. The rationale behind the ACTS effort, data related to ACTS effectiveness, and future research directions are also discussed. (Author)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

ED169916

R&D Utilization Report 78-1

U.S. DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY.

DEVELOPMENT AND EVALUATION OF AN ADAPTIVE COMPUTERIZED TRAINING SYSTEM (ACTS)

Bruce W. Knerr and Leon H. Nawrocki

Submitted by:

James D. Baker, Chief

EDUCATIONAL TECHNOLOGY AND TRAINING SIMULATION TECHNICAL AREA

Frank J. Harris, Acting Director
ORGANIZATIONS AND SYSTEMS
RESEARCH LABORATORY

Approved by:
[Signature]
TECHNICAL DIRECTOR

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES
5001 Eisenhower Avenue, Alexandria, Virginia 22333

Office, Deputy Chief of Staff for Personnel
Department of the Army

September 1978

Army Project Number
2Q161102B74B
2Q762717A764

Educational Concepts and
Evaluation

Approved for public release; distribution unlimited.

FOREWORD

The Educational Concepts and Evaluation Work Unit Area of the Army Research Institute for the Behavioral and Social Sciences (ARI) performs research and development in areas of educational technology with applicability to military training. Of special interest is research in the area of computer-based training systems. The development and implementation of such systems is seen as a means of reducing training time and costs by providing more highly individualized training than would be otherwise possible.

This report summarizes the research conducted in the development of one computer-based training system, the Adaptive Computerized Training System (ACTS). In order to accomplish this research, ARI's resources were augmented by contract with Perceptronics, Inc., an organization selected as having unique capabilities for research and development in this area.

The entire research work unit was initiated in FY 1974 in response to the requirements of EDT&E Project 2Q1672B74B, "Basic Research in the Behavioral and Social Sciences." The success of the initial effort resulted in a continuation of the research in response to the requirements of R&E Project 2Q762717A764, "Educational and Training Technology," also in response to the special requirements of the Product Manager, Computerized Training Systems, as expressed in Human Resources Needs Report 77-173 (Research Support for Computerized Training Systems) and 77-173 (R&E Support for Computerized Training Systems).


JOSEPE MEIDNER
Technical Director

TABLE OF CONTENTS

	Page
INTRODUCTION	1
The Need for Individualized Training	1
The Problem of Preparing Computer-Based Training Materials	2
The Promise of "Artificial Intelligence" Techniques	3
Improving Troubleshooting Training Procedures	4
Objectives	4
ACTS DESCRIPTION	5
ACTS Components	6
The Expected Utility Model	6
Developing ACTS Training Materials	9
Student Interactions	12
EVALUATION	15
ADAPTING TRAINING	20
FUTURE DIRECTIONS	25
REFERENCES	28
APPENDIX A. Bibliography of ACTS Publications	29
B. An Example of the Operation of the Expected Utility (EU) Model	31
GLOSSARY	35
DISTRIBUTION	41

LIST OF TABLES

Table 1. Measurement outcomes associated with circuit faults in the Heathkit EP-28 power supply	10
2. A sample set of student utilities	24
3. Sample decision rules and feedback statements	26

LIST OF FIGURES

	Page
Figure 1. ACTS circuit diagram	7
2. The ACTS student display as it appears at the start of a problem.	13
3. Student interaction sequence	14
4. Estimated relative utilities for normal and abnormal outcomes of two measurements (6 and 9) as a function of the number of decisions.	16
5. Student model predictive success during simulated student performance.	18
6. Estimated relative utility for key measurements as a function of fault problems	19
7. Student decision time as a function of training session.	21
8. Student decision efficiency (decisions/problem) as a function of training session	22
9. Student model predictive success as a function of fault problems	23

INTRODUCTION

In 1974 the Army Research Institute for the Behavioral and Social Sciences (ARI) initiated a program to develop an Adaptive Computerized Training System (ACTS). Since that time the program has produced three technical reports, three papers, computer software, and one systems user's manual.¹ However, no single report provides a comprehensive up-to-date summary of progress. This report is directed primarily toward a non-technical audience. Neither is there a report which places the program in the context of the military training problems which it is intended to solve. This report has been written to fill those needs. Therefore, the specific objectives of this report are:

- (a) To explain the rationale behind the ACTS research and development program;
- (b) To explain, in relatively non-technical terms, what the ACTS is supposed to accomplish and how it operates;
- (c) To provide a summary of the data evaluating ACTS effectiveness; and
- (d) To describe proposed future directions in ACTS research and development.

The overall objective of the ACTS program has been to develop and evaluate a new method for providing computer-based troubleshooting training.

There are a number of reasons for pursuing this objective, among which are the need for: cost-effective individualized training; simplified procedures for the production of computer-based training materials; evaluation of the applicability of advanced techniques from "artificial intelligence" research to Army training; and improved methods of training troubleshooting personnel. These areas will be examined in turn.

The Need for Individualized Training

The concept that training can be conducted more efficiently or effectively if adapted to characteristics (such as ability) or performance of the individual student is not a new one. Methods for providing such

¹A complete list of all publications produced as part of this effort is included as Appendix A, "Bibliography of ACTS Publications."

individualization were developed as early as the 1920's (Pressey, 1926) although not widely used. Not until the 1950's did individualization of training and education become a popular topic (Cronbach, 1957). The first approach to individualization to receive mass acceptance was Skinner's (1954) "teaching machine," and a variant, the linear programmed text. These methods permitted each student to proceed at his own pace, but instructional content remained the same for all students. The teaching machine varied the instructional sequence by requiring students to repeat segments to which they had previously responded incorrectly. In general, a very limited amount of individualization was not provided by the programmed text.

Crowder (1961) developed both a teaching machine and a text format which permitted branching (variation in instructional content and sequencing as a function of student responses) in addition to individual control of pace. Both the linear and branching approaches, under the general term Programmed Instruction (PI) have been widely adopted by the military and the civilian educational community. Many of the "new" advances in Army training, such as self-paced instruction and the Training Extension Course (TEC) audio-visual lessons, are applications of the principles of PI to non-paper-and-pencil media.

Both types of PI have been relatively successful. A branching PI text provides greater capability for individualization than does a linear one. However, the commonly used textbook format places practical limitations on the amount of individualization possible. The more branches included, the more cumbersome and difficult for the student to use the text.

The advent of real-time computers introduced the capability to provide greater individualization than is possible with a branching programmed text. Computer-Assisted Instruction (CAI), the use of the computer to provide instruction directly to the student, had its beginnings in 1959 (Kath, Anderson & Brainerd, 1959). CAI did not come into widespread use until the mid-1960's. Despite its tremendous potential to provide individualized instruction, CAI frequently has been a disappointment. Research evidence indicates that CAI is an effective, but not necessarily a practical, method for training and education.

The Problem of Preparing Computer-Based Training Materials

CAI systems frequently are referred to by observers as "automatic page turners." The appellation reflects the fact that the computer appears to do little more than present a new segment (page) of text after the student has made a response. While this is an oversimplification of the computer's function in many training applications, there is a grain of truth in it. Generally, capabilities of the computer for providing individualized training are underutilized. Several reasons would appear for this. A few reasons for this underutilization will be examined here.

The first reason is historical. As described above, CAI has been viewed primarily as a better method of presenting PI. Thus the practice of presenting instructional material in small segments has been transferred directly from PI to CAI. While its proponents claim that CAI is a flexible instructional medium, providing the capability for the use of a variety of instructional strategies, it has in fact been used extensively with only two of those strategies, drill and practice and tutorial, the latter being conceptually identical to the PI strategy.

A second reason is the difficulty, time, and cost required to prepare highly individualized CAI materials. As compared with the preparation of a PI text, the CAI author has the additional task of organizing (and often coding) the materials for machine execution. Moreover, the more options or branches within a lesson, the more complex the authoring process becomes. Increased complexity leads to increased preparation time, which in turn leads to increased cost of the lesson materials. When CAI is introduced in an operational setting, such as an Army school, the authors rarely have the luxury of providing individualization to the extent desirable. Lessons must be prepared within a fixed, and often very limited period of time. Individualization is often sacrificed for efficiency of production.

There are several possible ways to help the author in this situation. One is the development of on-line and off-line aids to assist the author in preparing lesson materials by performing the more tedious and time-consuming tasks. Another is through development of generalized instructional logic routines independent of instructional content. The author's task is then to insert the lesson content into a pre-existing framework. Unfortunately, the framework may not necessarily be the best one for any particular set of materials. Ideally, the instructional content, rather than the generalized logic, should determine the instructional sequence.

A third method for providing individualized lessons is the development of computer software systems which can themselves generate individualized training sequences through the use of an internal expert. This expert must be able to analyze student responses to determine the student's learning difficulties, and present the appropriate instruction to correct those difficulties. This last approach requires the use of "artificial intelligence" techniques.

The Promise of "Artificial Intelligence" Techniques

Artificial Intelligence (AI) techniques are algorithms (rules) which enable computers to exhibit "intelligent" behavior. Examples of intelligent behavior are understanding written English; playing chess, and learning (changing behavior as a result of experience). The field of AI developed in the 1960's as a tool for the study of human behavior. It was assumed that better understanding of complex human behavior could be obtained if the rules which enabled computers to produce the same behavior could be determined. By the 1970's it had become apparent that

the same techniques could be used to enhance CAI. Carbone's SCHOLAR system (1970) was the first to use an intelligent computer-based tutor, one that could both ask questions of the student and respond to unanticipated questions posed by the student.

The advantage of AI techniques is that they can provide highly individualized and flexible instruction without the necessity for programming the instructional logic separately for each lesson. The primary disadvantage is the extensive computer resources required to support CAI systems which use AI techniques. In the past, this has prevented the use of such systems outside a research environment. Fortunately, technological advances in computer hardware (particularly miniaturization) have resulted in a substantial reduction in the space requirements and initial cost previously associated with the computer capability necessary to support sophisticated AI.

Improving Troubleshooting Training Procedures

The Army's current approach to troubleshooting training is primarily hands-on and equipment specific. The student first is taught the sequential step-by-step procedures necessary to locate a malfunction in a particular item of equipment. Then he practices and is tested on the actual equipment itself. This approach has several advantages. It insures that the student has mastered certain prerequisite skills, such as the use of test equipment. This approach also teaches the student the physical layout of the equipment and the correspondence between the equipment and the circuit schematic diagram. Finally, it gives the student practice in disassembling and reassembling the equipment.

There are also several disadvantages. Since the training content is equipment specific procedures, rather than troubleshooting logic, there is little transfer of the skills acquired to similar or modified items of equipment. Also, a substantial amount of equipment, which otherwise could be used operationally, is required for training purposes. Instructors must spend large portions of their time inserting malfunctions into the equipment, rather than actually conducting training. Moreover, much student time is spent assembling, disassembling, and soldering the equipment, thus reducing the number of different equipment faults that they can experience during their training. The use of "intelligent" computer simulation offers one means of improving existing training procedures.

Objectives

When the ACTS program was initiated it was viewed as a solution to the problems described above. In summary, the overall objectives of the program have been:

- (a) To improve the individualization of training in CAI systems through ~~the~~ use of AI techniques;
- (b) To minimize the effort required by the lesson author;
- (c) To use relatively "basic" AI techniques which can be implemented on small-scale computer systems;
- (d) To evaluate the training and cost effectiveness of the system for electronic troubleshooting training.

ACTS DESCRIPTION¹

The ACTS program was initiated as a basic research effort in 1974. Work accomplished through January 1977 was performed by Perceptronics, Inc., under contract to and with guidance provided by ARI. The current version of the ACTS evolved gradually throughout this period. Previous reports (see Appendix A) described the ACTS as it existed when those reports were written, and, while historically accurate, do not always accurately reflect the current state of ACTS development. For example, the ACTS has been called the Computerized Diagnostic and Decision Training (CDDT) system and the Computerized Decision Training (CDT) system, as well as the ACTS, in different reports. This section describes the current version of the ACTS.

A brief overview of the ACTS will provide the background for the system description. The student's task in the ACTS training setting is to troubleshoot a complex electronic circuit by making various test measurements, replacing the malfunctioning part, and making final verification measurements. The entire process is simulated by the ACTS. Neither the actual circuit nor test equipment is required. The heart of the system is an adaptive computer program which "learns" the relative preference of the student for the various test measurements, compares this to those of an expert, and when complete, will provide feedback and adapt the instructional sequence to eliminate discrepancies between the student and the expert.

The ACTS is not being proposed as a complete troubleshooting training method. It has not been designed to train the student to use test equipment, assemble or disassemble the equipment, or provide him² with background information about the operation of the circuit. It is designed

¹A Glossary explaining the technical terms used in this section is included. Terms included in the Glossary are underlined when they first appear in the text.

²Both male and female personnel receive electronic troubleshooting training in the Army. The masculine gender in reference to those students is used only to avoid the awkward structure imposed by "his/her" wording.

to train the student in the decision-making aspect of the troubleshooting process.

ACTS Components

The ACTS consists of four major components: (a) the task model; (b) the expert model; (c) the student model; and (d) the instructional model.

Task Model. The task model is a simulation of the system (in this case an electronic circuit) on which the student is to be trained. The circuit currently being used is a modular version of the Heathkit IP 28 Power Supply.¹ A simplified schematic diagram of this circuit is shown in Figure 1. The power supply, when functioning properly, converts a 117-volt alternating current input (shown at the left) into a stable, low-voltage, low-amperage direct current output (shown at the right). As the diagram shows, the circuit consists of ten modules. Since the output of the circuit must be stable, even with variations in the input, there are a number of corrective feedback loops in the circuit which make the troubleshooting process more difficult.

Expert model. The second major component of the ACTS is a model of an expert troubleshooter. This is an Expected Utility (EU) model which predicts the expert's measurement choices as he troubleshoots the circuit. It is developed through on-line observation of the expert's troubleshooting behavior.

Student model. The student model, like the expert model, is an EU decision model which predicts the student's measurement choices. It is developed through on-line observation of the student's behavior as he solves troubleshooting problems on the ACTS.

Instructional model. The last major component is the instructional model. The function of the instructional model is to compare the expert and student models, determine discrepancies between the two, and to modify the instructional feedback and problem presentation sequence in order to reduce those discrepancies. Currently the instructional model can provide some adaptive feedback, but cannot modify the instructional sequence.

The Expected Utility Model

The uniqueness (and promise) of the ACTS lies in the use of the expert and student models. Since they are so important to ACTS

¹Commercial designations are used only for precision of description. Their use does not constitute endorsement by Department of the Army or the Army Research Institute.




Figure 1. ACTS Circuit Diagram (from Crooks, Kuppin, & Eredy, 1978).

operation, they will be examined in greater detail. While the two models serve different functions and use different data, their operation is identical.

Consider an expert troubleshooter who is given a defective IP-28 power supply and simply told "It doesn't work. Fix it." If he chooses to repair it, there are a limited number of actions that he can take. The troubleshooter can use the switches and meters on the front panel of the power supply to check any of the four outputs. He can take any of the 25 possible internal test measurements. Or this expert can replace any of the 10 circuit modules.

Each of these possible actions has associated with it a set of possible outcomes. For example, measuring the output voltage with the voltage setting in the high state could produce outcomes of normal, low, very low, or zero.¹ Module E, the current source, could be good or bad. If good, replacing module E would not correct the circuit malfunction. If bad, replacing the current source would correct the malfunction.² Each possible outcome has three properties. The first is the conditional probability of the occurrence of that outcome given that the appropriate action is selected and given the previous measurement outcomes. For example, given the previous measurement results, what is the probability that measuring the output voltage with the voltage setting in the high state will result in an outcome of zero? The second property is the utility of the outcome to the troubleshooter, i.e., what he gains or loses as a result of that outcome. Utility is subjective, but it should be related to the cost (in time or money) of taking that action. The third property is the gain in information that the outcome provides about the location of the fault. Information gain is a function of the number of faults in the circuit that would be eliminated if a particular result were obtained.

The expert and student models combine these three properties to obtain an "expected utility" for each possible action. The higher the expected utility of an action, the more desirable that action becomes. The properties are combined as follows:

$$EU_j = \sum_i^n \alpha_{ij} P_{ij} U_{ij} \quad (1)$$

¹When using the actual equipment, of course, the outcome would be a numeric value. The troubleshooter next would have to determine whether this value was high, medium, low, etc.

²Again, this is a simplification. The current source could be "bad" in a number of ways.

where

EU_j = the expected utility of action A_j

P_{ij} = the probability that outcome i , of a set of n outcomes, will occur if action A_j is selected

U_{ij} = the utility of outcome i of action A_j

α_{ij} = the information gain resulting from the occurrence of outcome i of action A_j .

In other words, the expected utility of any action is the sum, across all possible outcomes of that action, of the product of the probability, utility, and information gain of that outcome.¹

While it is assumed that human troubleshooters combine the properties in this fashion, and choose the action with the highest expected utility, this is not a necessary condition for ACTS operation. It is sufficient that the model predicts the actions of the human troubleshooter accurately.

An example of the operation of the expected utility model is shown at Appendix B.

Developing ACTS Training Materials

While the basic ACTS model is applicable to any diagnostic task, task-specific information must be provided before training can begin. The first step in obtaining this information is to prepare a table similar to Table 1. This table, developed for the IP-28 power supply, shows the outcome for each of the measurements for each of the possible circuit faults. This information forms the basis for the task model (circuit simulator) and is necessary for the development of the expert and student models.

The second step is to determine the probability of occurrence of each possible fault. The simplest procedure is to assume that all faults are equally likely to occur. This method is especially appropriate for new circuits for which no information about failure rates of the components is available. A more accurate representation of the on-the-job situation can be obtained by using probabilities which reflect the frequency of occurrence of each fault in the actual equipment. This information could be obtained from examination of maintenance records.

¹It should be noted that the ACTS EU model differs from the basic EU model described in the Glossary. Information, normally a component of utility, is treated separately in the ACTS.

Table 1

Measurement Outcomes Associated with Circuit
 Faults in the Heathkit IP-28 Power Supply
 (Blank spaces are Normal outcomes) (From Crooks, Kuppin, & Freedy, 1978)

MODULES	FAULTS	MEASUREMENTS																												
		P	Q	R	S	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
NORMAL	0																													
G. TRANSFORMER	1	L	H	L	H	F	F	F	Z	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	F	Z	F	Z	F	Z	Z		
	2	X				L	L	L	V	V	V	V	V	V	V	Z	Z	Z	V	Z										
	3	X				L	L	L	V	V	V	V	V	V	V	Z	Z	Z	V	Z										
	4	X				L	L	L	V	V	V	V	V	V	V	Z	Z	Z	V	Z										
	5	Z				F	F	F	Z	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	V	V	Z	Z	Z	Z	Z	Z	
B. DC POWER SOURCE	6	L	H	L	H	F	F	F	Z	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	V	V	Z	Z	Z	Z	Z		
	7					L	L	L	V	V	V	V	V	V	V	Z	Z	Z	V	Z										
E. CURRENT SOURCE	8	Z			H	F	F	F	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	V	V	Z	Z	Z	Z	Z	Z		
	9	Z			L	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
	10	Z			L	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
	11				Z	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
C. SERIES REGULATOR	12	L	H	L	H	F	F	F	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	V	V	Z	Z	Z	Z	Z			
	13				Z	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
	14	Z			L	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
D. CURRENT SENSE	15	Z	Z	L	H	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	Z	Z	V	V	V	V		
	16				L	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
	17				Z	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V	V		
F. CURRENT LIMITER	18				L	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	Z	Z	V	V	V	V		
	19	X	Z	L	H	F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		
K. REFERENCE DC SOURCE	20	X				F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		
	21	Z				F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		
L. VOLTAGE REFERENCE	22	Z				F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		
	23	Z				F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		
H. VOLTAGE LIMITER	24	Z				F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		
	25	L	H	L	H	F	F	F	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	V	V	Z	Z	V	V	V	V		
J. OUTPUT STAGE	26	L	H	L	H	F	F	F	Z	Z	Z	Z	F	F	Z	Z	Z	Z	Z	Z	V	V	Z	Z	V	V	Z	Z		
	27	X				F	F	F	Z	Z	Z	Z	V	V	V	V	V	V	V	V	Z	Z	V	V	V	V	V	V		

If the records are not available, estimates of the frequency of fault occurrence could be obtained from experts.

Given the table of faults and measurements and the probability of occurrence of each fault, the ACTS can calculate the conditional probability of occurrence and information gain of each outcome as shown in equation 1. When modules can contain more than one possible fault, the ACTS also can calculate the probability that particular module is bad. Equations for these calculations are shown in May, Crooks, and Freedy (1978).

Next, the costs for each measurement and module replacement must be determined. In general, measurement costs should reflect the amount of time required to take that measurement when using the actual equipment. The cost of replacing a module should reflect amount of time for replacement and cost of the new module. The costs are presented to the student in dollar amounts.

Programming the computer to display the circuit diagram is the last task before the expert model can be trained. This training process results in the utilities for the measurement outcomes as shown in equation 1. Referring once again to this equation, it can be seen that there are now two sets of unknowns remaining: the expected utilities for each possible action, and the utilities for each possible outcome. These unknowns are estimated by "tracking" an expert's troubleshooting behavior as he completes a series of problems on the ACTS. As he does this, he is presented with the updated probabilities of measurement outcomes. The values of the known model parameters (probability and information) are entered into the expert model before the expert starts, with the utilities set at some common arbitrary value (usually 100). The expert model chooses the action which has the highest expected utility. If the expert then chooses the same action, no changes in the model are made. However, if the action selected by the expert model differs from the action selected by the expert, the model utilities associated with the model choice are punished (decreased) and those associated with the expert choice are rewarded (increased). This process continues until the estimated utilities become stable. This will occur when the expert model is able to predict the choices of the expert accurately.

At this point, the expert is no longer needed. The expert model, having been "trained," replaces the human expert. Now the system is ready to begin training the student.¹ As did the expert, the student begins to troubleshoot. As he does this, he has access to the probabilities produced by the expert model.

¹Note that no mention has been made of techniques for modifying the instructional sequence or for providing feedback to the student. The methods for accomplishing this have yet to be determined, as will be discussed in a later section of this report.

As the student solves a series of problems, the student model, which functions in the same manner as does the expert model, learns the student's utilities. When the estimated student utilities begin to stabilize, feedback can be provided to the student.

Student Interactions

The student display, as it appears at the start of a problem, is shown in Figure 2. It has four areas: (a) the circuit area; (b) the main message area; (c) the considerations area; and (d) the legend/error-message area. The circuit area contains a diagram of the circuit, along with the measurements that the student can take, and the cost for each measurement. After a measurement has been made, the outcome is displayed instead of the cost. The considerations area is used to present the outcome probabilities to the student. The main message area informs the student of his options, and also provides instructions and assistance. The legend/error-message area provides a legend interpreting the codes by which the probabilities are displayed, and is also used to inform the student when he has taken an "illegal" action. All student inputs are accomplished through the use of a trackball and cursor.


Figure 3 presents a flow diagram of the student interactions with the ACTS during the training process. At the start of a problem, the student is told only that the circuit has a malfunction which he is to correct. Initially, the student can select one of five options. The first is to ask for HELP. As shown in Figure 3, this option is available to the student whenever he must select some action. HELP provides the student with two types of information from the expert model: (a) a list of the circuit modules that could be bad; and (b) the action that the expert would take next.

The second option that the student can select is to enter his choice of symptoms for consideration. Symptoms, which are measurements made on the final output of the circuit, are distinguished from the internal measurements. After the student chooses a symptom for consideration, the expert model's probabilities for the possible outcomes of a check of that symptom are shown. The student may choose to consider additional symptoms, select one of the symptoms considered, or choose none of the symptoms and return to the action selection choice point. HELP is also available. Selecting a symptom causes the outcome to be displayed, and then returns the student to the action selection choice point.

The third option that the student can select is to enter his choice of (internal) measurements for consideration. The sequence of interactions for this option is the same as for the previous option, entering a choice of symptoms for consideration.

CIRCUIT AREA

CONSIDERATIONS AREA



ACTIONS CONSIDERED

CHOOSE A SYMPTOM, MEASUREMENT, OR MODULE.
 REMEMBER TO LIST THE 4 ALTERNATIVES THAT YOU CONSIDER MOST IMPORTANT; THEN CHOOSE ONE FROM YOUR LIST.
 MOVE THE CURSOR TO THE POINT ON THE SCREEN THAT YOU WANT, THEN PRESS THE "ENTER" BUTTON.
 PRESS THE "ENTER" BUTTON ONLY WHEN THE READY INDICATOR IS ON.

HELP READY OK

MAIN MESSAGE AREA

LEGEND/ERROR-MESSAGE AREA

Figure 2. The ACTS Student Display As It Appears at the Start of a Problem






Figure 3. Student Interaction Sequence

The fourth option is to enter a choice of modules to be considered for replacement. After the student chooses a module for consideration, the expert model's probability that that module is defective is shown to the student. The student may consider additional modules, choose to replace a specific module, or return to the action selection choice point. Again, the expert model's probability is available. If the student chooses to replace a module, the replacement is simulated and the student is returned to the action selection choice point.

The student's final option is to declare the circuit OK. If the student is incorrect, he is given feedback to this effect and returned to the action selection choice point. If correct, the student is informed that he has corrected the circuit malfunction. At this point the utilities in the student model are updated and a performance summary is printed for the experimenter. If a block of problems has been completed, the student may be provided with feedback based on the student model utilities. Otherwise, a new problem is initiated.

EVALUATION

A one-year effort was required to develop the basic concepts for ACTS development and to produce the basic software to test those concepts. Since the basic software was completed at the beginning of 1975, the ACTS has undergone concurrent development and evaluation. Major items awaiting development are the utility-based feedback and the modification of the problem presentation sequence on the basis of student utilities. However, since evaluations have been conducted concurrently with the development process, there are sufficient data to show that the major requirements for successful ACTS operation have been met.

The most basic requirement is that the utility estimation algorithms (which estimate the expert and student utilities) operate correctly. This applies to both the expert model and the student model. Several tests of these algorithms have been conducted.

The first were conducted using pairs of selected measurements. An arbitrary set of utilities was chosen for the normal and abnormal outcomes of the two measurements under consideration.

A troubleshooting problem was initiated and an expert used the arbitrary utilities to calculate the expected utilities of those two measurements, according to the EU model described previously. The expert consistently chose the measurement with the greater expected utility. This process was repeated for several problems using different measurements. Figure 4 shows the change in the ACTS-derived utilities for the two outcomes (normal and abnormal) of two selected measurements. In this case, one of the utilities had an imposed value of two and the other three had an imposed value of one. This figure shows rapid stabilization (the utilities did not change after the second decision), and a




Figure 4. Estimated Relative Utilities For Normal and Abnormal Outcomes of Two Measurements (6 and 9) as a Function of the Number of Decisions (from May, Crooks, & Freedy, 1978).

rank ordering similar to that of the imposed values. These tests indicated that the utility training algorithm was operating correctly for pairs of isolated measurements.

A subsequent test was made by an expert who used a consistent overall decision strategy. Most of the available measurements were used. Following 14 problems (70 measurement-selection decisions), the utilities that had been adjusted by the ACTS stabilized at levels which corresponded to their ranking in the decision strategy used by the expert.

The same utilities were also inserted in a "simulated student" program to provide an additional test of the utility adjustment algorithm. Conditions were such that the student utilities were "known" and the decision-maker was completely consistent. The simulated student is a routine which troubleshoots the circuit using any set of utilities with which it may have been programmed. It always chooses the action with the highest expected utility. The resulting utilities produced by the utility estimation algorithm were in a similar rank order to the simulated student's utilities.

Since the expert model and the student model are, in essence, identical, it can be assumed that if the expert model functions properly, so does the student model. Nevertheless, this assumption was checked by conducting a test of the adaptive student model similar to that previously described for the expert model. The success of the student model in predicting the actions selected by the simulated student is shown in Figure 5. Accuracy increased rapidly during the first 80 decisions (approximately 18 problems), and perfect success was achieved after 210 decisions (45 problems).

The tests described above indicated that the adaptive utility-estimation algorithms operated correctly when the students or experts made decisions consistently. The next series of tests was conducted to determine how these algorithms performed when the decisionmakers were less than perfectly consistent.

The first test with a person not following a predetermined troubleshooting strategy used an expert electronics technician. He solved 30 troubleshooting problems using the ACTS while the utilities of the adaptive model were adjusted. In a post-session interview, the technician was asked why he selected specific measurements, and his estimates of the importance of those measurements. Those indicated as critically important in troubleshooting were identified as the key measurements. Figure 6 shows the adjustments to and stabilization of this expert's utilities for normal measurement outcomes of these key measurements. Their rank order is the same as his verbal ranking of their importance.

A subsequent study was conducted to examine the performance of experienced electronic technicians. Eight students who scored high on a




Figure 5. Student Model Predictive Success During Simulated Student Performance (from May, Crooks, & Freedy, 1978).




Figure 6. Estimated Relative Utility for Key Measurements as a Function of Fault Problems (from May, Crooks, & Freedy, 1978).

written electronics knowledge test were given 4 1/2 hours of experience on the ACTS, divided into three sessions. During the second of these sessions, they used the expert's estimates of action outcome probabilities as an aid in their selection of troubleshooting actions. During the first and third sessions, the probabilities were not provided to the students.

The students improved their decisionmaking speed throughout the three sessions. The mean and range of decision time performance are presented in Figure 7. Figure 8 shows student decision efficiency measured in decisions per problem. The figure also shows the students' dependence on the outcome probabilities provided by the expert. When the probabilities were withdrawn in the third session, the students' decision efficiency decreased.

The mean predictive success of the adaptive student model was also assessed during this test, as shown in Figure 9. During the second session, when the outcome probabilities were presented, the model predicted 75% of the students' choices correctly.

At this point in the evaluation process, the following conclusions can be made. First, the adaptive utility estimation algorithm can predict, after practice, the performance of a consistent decisionmaker. Second, the adaptive utility estimation algorithm accurately rank-orders the utilities of an expert technician. Third, the presentation of outcome probabilities improves both student performance and the predictive power of the student model. Finally, student performance improves with practice on the system in the absence of any feedback regarding utilities.

ADAPTING TRAINING

The ACTS is still not complete. The major difficulty has been the development of the mechanism (or mechanisms) by which the training is to be adapted to individual performance. Ideally such mechanisms should affect both the feedback that the student receives and the sequence in which problems are presented, and should be based on the student's utilities. This has turned out to be more complex than anticipated. One reason is the sheer volume of information provided by the utilities. There is one utility for each measurement outcome, and one for the replacement of each module. In the case of the IP-28 power supply, there are 96 different utilities. Table 2 shows a sample set of student utilities.

Thus far two methods of providing feedback to the student, one using student utilities, have been developed. Neither method alters the problem presentation sequence. The first, and simplest, method presents the expert model's action choice to the student after the student has made his selection and obtained the result. Student utilities do not play a part in this feedback sequence, but the expert model is required




Figure 7. Student Decision Time as a Function of Training Session (from May, Crooks, & Freedy, 1978).




Figure 8. Student Decision Efficiency (Decisions/Problem) as a Function of Training Session (from May, Crooks, & Freedy, 1978).




Figure 9. Student Model Predictive Success as a Function of Fault Problems (from May, Crooks, & Freedy, 1978).

Table 2

A Sample Set of Student Utilities

Measurement or Symptom Code	Utility for				Module Code	Utility for Replacement
	Normal Outcome	Non-normal Outcome 1	Non-normal Outcome 2	Non-normal Outcome 3		
P	100	100	100	100	B	100
Q	124	113	100	100	C	100
R	101	100	92		D	100
S	33	60	81		E	100
1	100	100	100		F	100
2	100	100	100		G	100
3	100	100	100		H	100
4	47	34	90		I	100
5	100	100	100		J	100
6	100	100	100		K	100
7	100	86	86		L	100
8	100	86	86			
9	100	100	100			
10	100	100	100			
11	100	100	100			
12	100	100	100			
13	100	100	100			
14	100	100	100			
15	100	100	100			
16	100	100				
17	100	113	113			
18	100	100	100			
19	100	100	100			
20	100	100	100			
21	100	100	100			
22	100	100				
23	100	100	100			
24	86	100	113			
25	100	100				

to make it practical. Determination of the "best" action at any point in the troubleshooting sequence must account for the previous actions taken and results obtained by the student. Between 4 and 12 sequential actions (decisions) are required to repair the circuit, and initially the student can choose from among 39 possible actions. To program all of the best choices using a logical branching technique would be, at best, a rather complex task.

The second method of providing feedback is based on comparisons among "key" student utilities. The key utilities are those for measurements identified by an expert, as being of critical importance in the fault isolation process. For the IP-28 power supply, the utilities for the outcomes for six measurements were considered to be key. Based on the relationships among these utilities, a set of six decision rules and feedback statements were developed. Samples are shown in Table 3. The decision rule for the first feedback statement should be read as follows: If any of the utilities for measurement 3 are less than any of the utilities for measurement 19, or if any of the utilities for measurement 3 are less than any of the utilities for measurement 11, present this feedback statement to the student.

The appropriate feedback statements initially are presented to the student after completing the 15th problem, with updated statements presented every 15 problems thereafter. The student can review them at any time.

FUTURE DIRECTIONS

At some point in the development and evaluation of any training system, the developers are faced with a dilemma: should continued efforts be made to improve the system, or should further development be stopped and the system evaluated in its current state? The first alternative offers a potentially better end product, but delays evaluation, and consequently, implementation. The second alternative offers a potentially less-than-optimal end product, but earlier evaluation and implementation.

Such is the dilemma currently facing the ACTS. On one hand, there are a number of questions remaining to be resolved and improvements to be made, particularly with regard to providing feedback to the student and altering the problem presentation sequence. On the other hand, the ACTS appears to have the potential, as it currently exists, to improve Army troubleshooting training. For these reasons, future ACTS research and development will follow two initially divergent, but converging, paths.

One path will be directed toward evaluating the training and cost effectiveness of the ACTS in an ongoing course of instruction at an Army school. As a prerequisite to this, laboratory evaluations will be conducted to determine: (a) the transfer of ACTS training to the actual equipment; (b) the relative proportions of training with the ACTS and

Table 3

Sample Decision Rules and Feedback Statements

Rule	Feedback
3<19 or 3<11	Although measurement 3 is located at a good point to isolate the power input modules, it is expensive. Use this measurement after you have eliminated most other possibilities. Measurement 3 should be used when the probability of a normal outcome is rather high but not certain (a range of 60% to 80%).
2<11 or 3<11 or 4<11	A good first step in checking the operation of current and voltage feedback loops is to check the output of the series regulator. This should be done with the circuit operating at full output since this fully exercises the circuit functions. Therefore, measurement 9 or 11 should be used even if there is a low probability of a normal outcome. Use measurement 11 since it is much cheaper than 9.

training with the actual equipment required for optimal training effectiveness; and (c) the effects of variations in the problem presentation sequence on student performance. Changes made in ACTS software will, in general, be limited to those required to install the system at an Army school. Guidelines for the use of the ACTS will also be prepared. Finally, an evaluation will be conducted at the school.

Concurrently, research will be conducted as a basis for the development of a second-generation ACTS. Research topics will include the extent to which expert utilities agree, the use of student utilities as a criterion for stopping training and evaluating student performance, and the development of diagnostic measures based on student utilities.

REFERENCES

- Carbonell, J. R. AI in CAI: An Artificial Intelligence Approach to Computer Assisted Instruction. IEEE Transactions on Man-Machine Systems, MMS-11, New York, December 1970.
- Cronbach, L. J. The Two Disciplines of Scientific Psychology. American Psychologist, 1957, 12, 671-684.
- Crooks, W. H., Kuppin, M. A., & Freedy, A. Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction: Adaptive Computerized Training System (ACTS) (ARI Technical Report 78-A6). Alexandria, VA: US Army Research Institute, 1978.
- Crowder, N. A. Automatic Tutoring by Intrinsic Programming. In A. A. Lumsdaine & R. Glaser (Eds.), Teaching Machines and Programmed Learning, National Educational Association, Washington, D.C., 1960.
- May, D. M., Crooks, W. H., and Freedy, A. Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction: Experimental Studies (ARI Technical Report TR-78-A4). Alexandria, VA: US Army Research Institute, February 1978. (NTIS No. AD A050887)
- Pressey, S. L. A Simple Apparatus Which Gives Tests And Scores--And Teaches. School and Society, March 20, 1926, 23.
- Rath, G. J., Anderson, N. S., and Brainerd, R. C. The IBM Research Center Teaching Machine Project. In E. Galanter (Ed.), Automatic Teaching; The State of the Art. New York: John Wiley & Sons, 1959.
- Skinner, B. F. The Science of Learning and the Art of Teaching. Harvard Educational Review, 1954, 24, 86-97.

APPENDIX A: BIBLIOGRAPHY OF ACTS PUBLICATIONS

Crooks, W. H., Kuppin, M. A., & Freedy, A. Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction: Adaptive Computerized Training System (ACTS) (ARI Technical Report 78-A6). Alexandria, VA: US Army Research Institute, 1978.

Freedy, A., and Crooks, W. H. Use of an Adaptive Decision Model in Computer-Assisted Training of Electronic Troubleshooting, (In Proceedings of the Conference on New Concepts in Maintenance Training Research. Orlando, Florida: April 1975.)

Knerr, B. W. Artificial Intelligence Techniques in Electronic Troubleshooting Training: Development and Evaluation. In Training: Technology to Policy, Proceedings of the 1976 Symposium of the Potomac Chapter, Human Factors Society, (ARI Special Publication P-77-6). Alexandria, VA: US Army Research Institute, December 1977. (NTIS No. AD-A049629)

Knerr, B. W. Problems of Performance Measurement in Computer-Based Simulation. Paper presented at the Computer Adaptive Testing '77 Conference, Minneapolis, Minn., July 1977.

Kuppin, M. A. Adaptive Computerized Training System (ACTS): User's Manual (Perceptronics System Document 1028-76-1). Woodland Hills, Calif.: Perceptronics, Inc., December 1976.

May, D. M., Crooks, W. H., & Freedy, A. Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction: Experimental Studies (ARI Technical Report TR-78-A4). Alexandria, VA: US Army Research Institute, February 1978. (NTIS No. AD-A050887)

May, D. M., Crooks, W. H., Purcell, D. D., Lucaccini, L. F., Freedy, A., & Weltman, G. Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction (ARI Technical Report TR-77-A26). Alexandria, VA: US Army Research Institute, October 1977.

APPENDIX B: AN EXAMPLE OF THE OPERATION OF THE EXPECTED UTILITY (EU) MODEL

There is a very simple circuit which can fail in only one of four possible ways. These possible faults are labeled A, B, C, and D. There are also only four possible test measurements that can be taken. These measurements are identified by the numbers 1, 2, 3, and 4, and have costs associated with them of \$1, \$2, \$3, and \$4, respectively. Each measurement can have one of three possible outcomes, Normal (N), Low (L), or High (H). The following table shows the relationships between the faults and the possible measurement-outcomes.

Faults	Measurement			
	1	2	3	4
OK	N	N	N	N
A	L	L	H	H
B	L	N	H	L
C	N	N	H	L
D	H	H	H	H

All faults are equally likely to occur. When the troubleshooter is given the circuit for diagnosis, he knows that the circuit is not "OK." Before any measurements are taken, the probability of occurrence of each of the measurement outcomes is as follows:

Measurement	Outcome		
	L	N	H
1	.50	.25	.25
2	.25	.50	.25
3	.00	.00	1.00
4	.50	.00	.50

The amount of information resulting from each outcome is:

Measurement	Outcome		
	L	N	H
1	.50	.50	.50
2	.50	.50	.50
3	.00	.00	.00
4	.50	.00	.50

Assume that the utilities for all possible outcomes (L, N, and H) of any measurement are the same, and that that value is equal to the reciprocal of the cost of taking that measurement. The utilities for the measurement outcomes are then:

Measurement	Outcome		
	L	N	H
1	1.00	1.00	1.00
2	.50	.50	.50
3	.33	.33	.33
4	.25	.25	.25

The expected utility for each measurement is calculated as the product of the probability, utility, and information for each outcome, summed across all outcomes for any measurement.

Measurement	Outcome			Expected Utility
	L	N	H	
1	.25	.125	.125	.50
2	.0625	.125	.0625	.25
3	.00	.00	.00	.00
4	.0625	.00	.0625	.125

Measurement 1 would be chosen by the EU model because it has the highest expected utility.

Now assume that measurement 1 is taken and an outcome of L is obtained. This eliminates faults C and D from consideration. The table showing the relationships between faults and outcomes can be reduced to the following:

Fault	Measurement		
	2	3	4
A	L	H	H
B	N	H	L

The probabilities of measurement outcomes are now:

Measurement	Outcome		
	L	N	H
2	.50	.50	.00
3	.00	.00	1.00
4	.50	.00	.50

The information resulting from each outcome is:

Measurement	Outcome		
	L	N	H
2	.50	.50	.00
3	.00	.00	.00
4	.50	.00	.50

The utilities for each outcome remain unchanged. The expected utilities for each measurement then become:

Measurement	Outcome			Expected
	L	N	H	Utility
2	.125	.125	.00	.25
3	.00	.00	.00	.00
4	.0625	.00	.0625	.125

The expected utility model in this case would choose measurement 2.

GLOSSARY

This glossary contains an alphabetical list of 11 technical terms used in this report. In parentheses immediately following the term itself is a list of other terms, also in the glossary, with which the reader should be familiar before reading that definition.

ALGORITHM -- A rule for performing a computation or solving a problem. For example: "To find the sum of a set of consecutive numbers, one through N , multiply the largest number (N) by one plus the largest number ($N + 1$) and divide the result by two."

CONDITIONAL PROBABILITY (PROBABILITY) -- The probability that a particular outcome (A) will be obtained if some other outcome (B) has previously been obtained, or, in other words, the probability of outcome A conditional upon B. Some outcomes are independent: the occurrence of outcome B has no effect on the probability of outcome A. For example, the outcome obtained on one throw of a die has no effect on the probability of obtaining a particular outcome on another throw of the same die. Conditional probability is relevant only to non-independent events. Assume that two consecutive throws of a die are to be made, and that the favorable outcome is a total of six on the two throws. Before the first throw is made, there are 36 possible outcomes, five of which (1-5, 2-4, 3-3, 4-2, and 5-1) will produce a total of six. The probability of obtaining a total of six is therefore $5/36$, or approximately 0.14. Now assume that the first throw is made and a four is obtained. There are now six possible totals that could be obtained (5, 6, 7, 8, 9, and 10), only one of which is six. The conditional probability of obtaining a total of six given that the outcome of the first toss was a four is therefore $1/6$, or approximately 0.17.

COST (VALUE) -- Negative value.

DIAGNOSTIC TASK -- Any task which requires an individual to determine the cause or nature of a problem or situation. Electronic troubleshooting, mechanical maintenance, and medical diagnosis are common types of diagnostic tasks.

EXPECTED UTILITY (EU) (PROBABILITY, VALUE, COST, UTILITY) -- A numerical expression of the anticipated subjective worth to an individual of taking an action, when the outcome that will result from taking that action is uncertain. Mathematically, if there are N possible outcomes, each having a probability of occurrence P_N , and a utility U_N , then

$$EU = P_1U_1 + P_2U_2 + P_3U_3 + \dots + P_NU_N$$

or

$$EU = \sum_{i=1}^N P_i U_i$$

As an illustration, consider the following situation. You have just arrived in a distant city for an important meeting the following day. In your hasty departure, you forget to pack your raincoat. Weather reports indicate a probability of 0.50 that it will be fair tomorrow, a probability of .25 that there will be light rain, and a probability of 0.25 that there will be very heavy rain. You definitely will be exposed to the weather, and it is very important to you that you arrive at your meeting in a dry condition. You enter a clothing store which offers two types of raincoats: inexpensive, which will protect you against a light rain, but offers little protection against a heavy rain; and expensive, which will protect you against light or heavy rain. You cannot wait until tomorrow to decide what to do. You carefully consider your three possible actions (buy nothing, buy the inexpensive raincoat, and buy the expensive raincoat) and the three possible weather conditions (fair, light rain, and heavy rain), and determine your utility for each action under each weather condition. The results are as follows:

ACTION	WEATHER CONDITION		
	fair (P=.50)	light rain (P=.25)	heavy rain (P=.25)
buy nothing	0	-30	-100
buy inexpensive raincoat	-15	0	-30
buy expensive raincoat	-75	-65	0

If you buy the appropriate raincoat for the weather, you "break even." Buying too little protection is unfavorable because your appearance will be degraded and your clothing will be damaged. Buying too much protection is unfavorable because you will have spent money unnecessarily. Which action has the highest expected utility? Using the formula shown above:

$$\begin{aligned} \text{EU (buy nothing)} &= (0.50)(0) + (0.25)(-30) + (0.25)(-100) \\ &= -32.00 \end{aligned}$$

$$\begin{aligned} \text{EU (buy inexpensive raincoat)} &= (0.50)(-15) + (0.25)(0) \\ &\quad + (0.25)(-30) \\ &= -8.25 \end{aligned}$$

$$\begin{aligned} \text{EU (buy expensive raincoat)} &= (0.50)(-75) + (0.25)(-65) \\ &\quad + (0.25)(0) \\ &= -53.75 \end{aligned}$$

Buying the inexpensive raincoat is the action which has the highest expected utility.

EXPECTED UTILITY MODEL (VALUE, COST, UTILITY, PROBABILITY, EXPECTED UTILITY) -- A model of human decisionmaking which assumes that, when individuals must select one of a number of alternative actions, they select the action with the highest expected utility.

INFORMATION (PROBABILITY) -- Information can be loosely defined as anything that reduces uncertainty about an outcome. Information is measured in binary digits, or bits. One bit is the amount of information that will eliminate one-half of the alternative outcomes, assuming that all outcomes are equally likely. Being told that the outcome of a coin toss was a head, or that the outcome of a roll of a die was either one, two, or three each convey one bit of information. Knowing which of four equally likely outcomes has occurred provides two bits of information, and which of eight, three bits. If the alternatives are not equally likely, the amount of information in an N-alternative situation is

$$-P_1 \log_2 1/P_1 - P_2 \log_2 1/P_2 - P_3 \log_2 1/P_3 - \dots - P_N \log_2 1/P_N$$

or

$$\sum_{i=1}^N p_i \log_2 1/p_i$$

where p_i is the probability of occurrence of outcome i .

PARAMETER -- A variable term in a mathematical function which determines the specific form of the function, but does not affect its general nature. For example, the function

$$y = ax + b$$

describes a straight line. The parameters a and b determine the slope and intercept, respectively, of that line. The function remains that for a straight line no matter what specific values may be substituted for a and b .

PROBABILITY -- A numerical expression of the likelihood of obtaining a particular outcome, usually called the "favorable" outcome. The numerical value of a probability can range from zero (the favorable outcome certainly will not be obtained) to one (the favorable outcome certainly will be obtained). Probabilities may be either objective, based on accepted rules for their calculation, or subjective, based on an individual's opinion or belief. Objective probabilities can be calculated in several ways. If it can be assumed that all possible outcomes are equally likely, the probability of occurrence of the favorable outcome is the number of possible favorable outcomes divided by the total number of possible outcomes. For example, the probability of obtaining the outcome "head" when a coin is tossed is $1/2$, or 0.5 ; the probability of obtaining the outcome "3" when a die is rolled is $1/6$, or approximately 0.17 ; the probability of obtaining the outcome "queen of hearts" when drawing a single card from a deck of playing cards is $1/52$, or approximately 0.02 ; and the probability of obtaining the outcome "any heart" when drawing a single card from a deck of playing cards is $13/52$, or 0.25 . The probability of occurrence of the favorable outcome can also be defined as the relative frequency of occurrence of that outcome that has been observed in the past. For example, assume that a jar contains a very large (infinite) number of marbles. One hundred marbles have been drawn from this jar. Of these, 50 were red, 25 were blue, and 25 were white. The probability of drawing a red marble on the next draw is $50/100$, or 0.50 ; that of drawing a blue marble $25/100$, or 0.25 ; and that of drawing a white marble $25/100$, or 0.25 . Subjective probabilities may be obtained for outcomes for which objective probabilities cannot be obtained (for example, the probability that the Tampa Bay Buccaneers will win Super

Bowl XX), as well as outcomes for which objective probabilities can be obtained. In the latter case, subjective and objective probabilities for the same outcome will not necessarily be identical. In both cases, subjective probabilities for the same outcome may vary from individual to individual.

UTILITY (VALUE, COST) -- A numerical expression of the subjective worth of an outcome to an individual. Utility is positive if the outcome is favorable, negative if the outcome is unfavorable, and zero if the outcome is neither favorable nor unfavorable. Utilities for the same outcome can vary from individual to individual, even for outcomes such as monetary gains or losses.

VALUE -- A numerical expression of the objective worth of an outcome. The value of a \$1.00 bet on a football game is +\$1.00 if the outcome is favorable (you win) and -\$1.00 if the outcome is unfavorable (you lose). The term cost is frequently used for negative value.