AD/A-003 015

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES: AN INITIAL STUDY IN DYNAMIC UTILITY CONVERGENCE AND DECISION AIDING

Richard L. Weisbrod, et al

Perceptronics, Incorporated

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

29 November 1974

DISTRIBUTED BY:

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

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER PTR-1016-74-11	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER AD/A-003015
4. TITLE (and Subtitle) ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES: AN INITIAL STUDY OF DYNAMIC UTILITY CONVERGENCE AND DECISION AIDING		5. TYPE OF REPORT & PERIOD COVERED Semi-Annual Technical Report 4-1-74 to 9-30-74 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Richard L. Weisbrod, Kent B. Davis, Amos Freedy, and Gershon Weltman		8. CONTRACT OR GRANT NUMBER(s) N00014-73-C-0286
9. PERFORMING ORGANIZATION NAME AND ADDRESS Perceptronics, Inc. 6271 Variel Avenue Woodland Hills, California 91364		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Work Unit NR 196-128 ARPA Order# 2347
11. CONTROLLING OFFICE NAME AND ADDRESS		12. REPORT DATE November 29, 1974 13. NUMBER OF PAGES 59
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) Engineering Psychology Programs, Code 455 Office of Naval Research 800 North Quincy Street Arlington, Virginia 22217		15. SECURITY CLASS. (of this report) Unclassified 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

None

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

- Command
- Control
- Computer Aided Decision
- Adaptive Learning Systems
- Man-Machine Interaction
- Man-Computer Interaction
- Utility Estimation

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

One central goal of the decision theory is to provide a rational basis for decision making. The ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System is designed to aid the decision maker (DM) in performing dynamic decision tasks. The ADDAM system provides real-time dynamic assessments of multiple utilities as the DM performs a dynamic decision task. ADDAM continuously tracks the DM's decision responses and uses adaptive pattern classification techniques to learn his utilities for their outcomes. These

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Utilities are then used to provide decision aiding of various forms, including recommendation of maximum expected utility decisions.

The results of an initial experiment to investigate the effectiveness of the adaptive decision modeling system and the DM's acceptance of the model as a normative basis for decision making are reported herein. It was found that the multiple dynamic utility estimates converged quickly to stable values. The model was sensitive to individual differences in decision strategy and was highly accurate in predicting the DM's decisions. Decision aiding, coupled with knowledge of the adaptive nature of the aiding, significantly increased decision making consistency in terms of the normative decision model. Aiding also resulted in an increase in decision speed throughput, though speed was not a performance criteria in the experiment.

PTR-1016-74-11
November 29, 1974

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES: AN INITIAL STUDY IN DYNAMIC UTILITY CONVERGENCE AND DECISION AIDING

RICHARD L. WEISBROD, KENT B. DAVIS,
AMOS FREEDY, GERSHON WELTMAN

This research was supported by the advanced Research Projects Agency of the Department of Defense and was monitored by Engineering Psychology Programs, ONR, under Contract No. N0014-73-C-0286 (NR-196-128)

PERCEPTRONICS

6271 VARIEL AVENUE • WOODLAND HILLS • CALIFORNIA 91364 • PHONE (213) 884-7470

i a

TABLE OF CONTENTS

		Page
1.	SUMMARY	1-1
	1.1 Purpose	1-1
	1.2 Methodology	1-1
	1.3 Results	1-2
	1.4 Future Work	1-4
	1.5 Organization of Report	1-4
2.	ADDAM SYSTEM OVERVIEW	
	2.1 The Decision Task	2-1
	2.2 The ADDAM System	2-8
3.	PILOT STUDY	3-1
	3.1 Objectives	3-1
	3.2 Description	3-1
	3.3 Results and Discussion	3-4
4.	INITIAL STUDY OF CONVERGENCE AND DECISION AIDING	4-1
	4.1 Objectives	4-1
	4.2 Hypotheses	4-2
	4.3 Experimental Design	4-2
	4.4 Subjects and Procedures	4-3
	4.5 Experimental Task	4-4
5.	RESULTS AND DISCUSSION	5-1
	5.1 Utility Measurement	5-1
	5.2 Decision Aiding	5-9
6.	REFERENCES	6-1

LIST OF FIGURES

Figure		Page
2-1	Decision Task Sequence	2-5
2-2	ADDAM System Task Work Station	2-6
2-3	Operator Information Display	2-7
2-4	Intelligence Analysis Report	2-9
2-5	ADDAM Functional Organization	2-10
2-6	Utility Matrix	2-13
2-7	Schematic Representation of Dynamic Utility Estimator	2-15
3-1	Estimated Utilities for Information from Trawler Sensors	3-5
3-2	Estimated Utilities for Information from Something Sensors	3-6
3-3	Estimated Utilities for Information from Everything Sensors	3-7
5-1	Estimated Utilities for Information from Iceberg Sensors	5-3
5-2	Utility for Information that an Iceberg is Present, from Several Types of Sensors	5-5
5-3	Utility for Information that an Iceberg is Not Present, from Several Types of Sensors	5-6
5-4	Indifference Point for Iceberg and Something Sensors	5-7
5-5	Variation in Indifference Point for Iceberg and Something Sensors Based on Utility Estimates (For Third Session)	5-8
5-6	Model Derived Expected Utility Indifference Point vs. Subjects Stated Indifference Point for Iceberg and Something Sensors	5-11

LIST OF TABLES

Table		Page
2-1	Sensor Properties	2-3
3-1	Strategy Rules SR 1	3-3
3-2	Strategy Rules SR 2	3-3
4-1	Experimental Design	4-3
4-2	Generalized Strategy Rules	4-5
5-1	Model Predictions of Sensors Actually Deployed by Control Subjects During the Fourth Session	5-2
5-2	Predicted vs. Stated Indifference Point Between Iceberg and Something Sensors	5-10
5-3	Utility Variation During Fourth Session (N=9)	5-13
5-4	Performance Rate During Fourth Session (N=9)	5-15
5-5	Normalized Utility Variation During Fourth Session (N=9)	5-16

1. SUMMARY

1.1 Purpose

One central goal of decision theory is to provide a rational basis for decision making. A normative model frequently used in situations involving risk is the expected utility model, a linear model based on the decision maker's (DM) utilities for the possible outcomes of his decisions and the probabilities of their occurrence. The utility assessments used in such decision models are usually derived by what might be termed static assessment techniques (Kneppreth, Gustafson, Johnson, and Leifer, 1973). These techniques may not be adequate for dynamic decision tasks, particularly where multi-attribute utilities are involved.

The ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System represents the application of an adaptive technique for dynamic utility assessment to the problem of decision aiding in a dynamic decision environment (Freedy, Weisbrod, Davis, May, and Weltman, 1974). This report presents the results of an initial experiment conducted to investigate the effectiveness of an adaptive model of DM behavior and the DM's acceptance of this model as a normative basis for decision making. Also discussed are the results of a pilot study conducted prior to the initial experiment.

The work reported herein was done as part of a research project whose overall goals are to implement promising techniques for adaptive modeling of human decision making; to explore factors which influence effective monitoring, aiding, and automating of dynamic decision processes; and to establish guidelines for the application of adaptive decision systems.

1.2 Methodology

The basic methodology was to have the experimental subjects perform a dynamic decision task under controlled conditions. The task required the subjects to track a fishing fleet in a simulated real-world environment by deploying sensors of varying response specificity, reliability, and cost. ADDAM continuously tracked the operator's (DM's) decision responses in real time and learned his decision strategy by performing the task in parallel with him. In effect, the decision maker "showed" the computer how to optimize on his own terms. ADDAM then continued the process and served as a normative model for aiding the DM.

The initial experiment was a one-way factorial design using nine subjects in three groups. Each subject had four 1-1/2 hour sessions. The first three sessions were training sessions, during which the subjects learned the task and ADDAM learned the subject's behavior. After the third session the three groups were given differential treatment. Group I, the control group, received no decision aiding. Group II received decision aiding in the form of sensor deployment recommendations derived from the adaptive decision model. This group was not told how the recommendations were derived. Group III subjects received decision aiding and were given a brief explanation of how the recommendations were derived from their own behavior. Group II subjects thus had a low degree of knowledge of the nature of the aiding while Group III subjects had a high degree of knowledge.

The measures of interest were ADDAM's estimates of the operator's utilities; the Utility Matrix Difference (UMD) score, a measure of the variability of utility estimates; the accuracy of decision predictions; and the total number of decisions made by the operator.

1.3 Results

ADDAM was found to be highly effective in tracking and predicting the operator's decision behavior. The estimates of multiple dynamic utilities converged quickly to stable and distinct values and the model was found to be very accurate in predicting the operator's decisions. During the fourth session, at least 95% of the unaided operator's (Group I) decisions were accurately predicted.

The adaptive decision model was found to be sensitive to individual differences in decision strategies. One such individual difference is the probability at which a subject is indifferent between two different sensors which report more or less equivalent information. Indifference probabilities computed from the dynamic utility estimates were compared with those elicited directly from the subjects. The correlation coefficient of 0.82 was significant at the 0.01 level.

Decision aiding presented to subjects who had knowledge of the adaptive nature of the aiding (Group III subjects) resulted in a higher degree of consistency with the normative decision model. This was indicated by a significant reduction in the variability of ADDAM's estimates of the operator's utilities. On the other hand, aiding without knowledge of the nature of the aiding appeared to accentuate individual differences in behavior.

Decision aiding appeared to improve the decision throughput of the subjects by allowing them to place sensors more quickly and by reducing the amount of vacilation near the indifference points. While this improvement had only marginal statistical significance, decision speed was not stressed as an important performance criteria. Undoubtedly, placing emphasis on speed will accentuate the improvement in throughput. Partially or completely automating the decision process (subject to operator override) will also increase the rate at which decisions are made. Comparisons of the indoctrinated and unindoctrinated aiding groups (Groups III and II) suggests that the operator's knowledge that the automated decisions are essentially his own will make this form of rapid decision making both acceptable and effective.

1.4 Future Work

A number of changes to the ADDAM system are currently being made to facilitate additional experimentation with the system. The utility estimator is being modified to improve the rate at which it is trained, and new sensors with a greater degree of decision strategy flexibility are being implemented. Features which give the operator his payoff score and, at the experimenter's option, allow him to get real-world validation of his sensor outputs and status decisions are being incorporated into the system. Also, a much more extensive facility for monitoring system performance and recording experimental data is being implemented.

Additional experiments are being planned to further study the validity and sensitivity of the model and the effects of decision aiding. How well does the model predict behavior in a wider range of decision strategies? How well does it respond to changes in the operator's decision strategies? How long does it take to respond, and how large must these changes be before they are detectable? What kind of behavior is the model best able to predict? Least able to predict? Does aiding improve operator performance? Quality of decisions? Consistency? Speed? What factors influence the effectiveness of aiding and what forms of aiding produce optimal results? These are a number of the questions toward which future research will be directed.

1.5 Organization of Report

Chapter 2 gives a brief overview of the ADDAM system and its operation. It also describes the decision task. Chapter 3 describes the pilot study conducted with the system and the results. The initial experimental study and its results are described in Chapters 4 and 5, respectively.

2. ADDAM SYSTEM OVERVIEW

The ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System is a flexible vehicle for research on dynamic decision theory, adaptive decision models, dynamic utility estimation, and man/computer decision making. ADDAM combines a system for simulating a dynamic decision task (Freedy, May, Weisbrod, Weltman, 1974) with an adaptive decision model based on dynamic utility estimation (Freedy, Weisbrod, Davis, May, and Weltman, 1974). It also includes mechanisms for man/computer interaction and decision aiding.

2.1 The Decision Task

The decision task is to gather intelligence about a dynamically varying hierarchical organization -- a simulated fishing fleet moving in an expanse of ocean -- and to report its status. To gather this information the operator deploys a variety of sensors with different response specificities, reliabilities, and costs. He then integrates the information he receives into a status report and continues to gather information about the objects in the environment as they move about. The decisions currently under study are the operator's sensor deployment decisions.

Environment. The simulated environment is a homogeneous expanse of ocean which has been divided into a 25 square (five by five) spatial grid. This grid is referred to as the board. For the experiments reported herein, the fishing fleet consists of a trawler which moves from square to square and periodically deploys its nets. Also present is an iceberg which moves around the board.

Several time-varying environmental conditions affect the behavior of the trawler, nets, and iceberg. These conditions include time of day (day or night), weather (clear or stormy), and phase of moon. Proximity to the iceberg also affects the trawler's behavior. The operator is not made explicitly aware of the environmental conditions or their effect on the behavior of the objects. However, the environmental conditions are known to the simulated intelligence analysis experts (described below) and are implicitly reflected in the intelligence report which is presented to the operator.

Each object on the board has several characteristics: type (iceberg, trawler, trawler with nets deployed), location, and heading (north, east, south, west, null). The operator cannot observe the characteristics of the objects directly, but can only infer them from the outputs of the sensors he deploys.

Sensors. The operator has several different kinds of sensors available for detecting objects in the environment. These sensors vary in object sensitivity, response specificity, false alarm rate, and cost. Object sensitivity refers to the kind of objects a sensor can detect and response specificity refers to the kind of response the sensor can give to detected objects. An unlimited number of sensors of each type are available to the operator, but he can deploy only one sensor per square on the board and must pay a cost for each sensor he deploys. A deployed sensor responds to objects within its immediate square only. The sensors never fail to respond to an object, but sometimes they report false alarms.

Six different kinds of sensors were available to the operators during the pilot study (Chapter 3) and the preliminary experiment (Chapter 4). These sensors included two kinds of trawler sensors, a net sensor, an iceberg sensor, an "everything" sensor, and a "something" sensor. Both the everything and something sensors responded to every type of object, but the something sensor, unlike any of the other sensors, does not identify what kind of object it detected. The two trawler sensors differed from each other in cost and error rate. The properties of the sensors are defined in Table 2-1. During

TABLE 2-1

Sensor Properties

Type	Object Sensitivity	Response Specificity	False Alarm Rate	Cost
T1	Trawler	Trawler	0.10	4.0
T2	Trawler	Trawler	0.20	2.0
N	Trawler with Net	Net	0.20	2.0
I	Iceberg	Iceberg	0.20	2.0
E(Everything)	Trawler/Net/Iceberg	Trawler/Net/Iceberg	0.05	8.0
S(Something)	Trawler/Net/Iceberg	Positive/Negative	0.30	1.0

the pilot study the cost of a something sensor was 2.00. This, subsequently, was changed to make it more consistent with the false alarm rate.

<u>Decision Task Sequence</u>. The decision task sequence (Figure 2-1) begins when the operator deploys his sensors. Once he has finished deploying his sensors, the sensor outputs are displayed in front of him. Some of the deployed sensors may give positive responses while others may not. On the basis of the sensor responses, knowledge of sensor behavior, previous sensor responses, etc., the operator reports what he believes is the status of the environment. The operator reports the object type, location, and heading. This information, plus the environmental conditions are used to generate an intelligence analysis report based on an "expert's" assessment of the situation. The intelligence report is displayed to the operator on a teletype. It is also used by the adaptive decision model. Decision aiding information derived from the decision model is then displayed to the operator (depending on experimental conditions) and the cycle begins anew with the deployment of new sensors.

<u>Work Station</u>. While performing the decision task, the operator sits at the work station illustrated in Figure 2-2. This work station consists of a graphics display terminal with keyboard input and a teletypewriter. The board, sensor deployments and outputs, status report, and decision aiding are displayed on the graphics terminal (Figure 2-3) at appropriate times during the decision task cycle. All operator inputs, i.e., sensor decisions and status reports, are transmitted to ADDAM via the keyboard. The teletype is used to print out the intelligence report. Experimental data, in coded format, is also printed on the teletype.

<u>Intelligence Analysis Report</u>. The intelligence analysis report is generated by a simulated "intelligence analysis expert". The report gives the probabilities of finding icebergs, trawlers, nets, and something at each

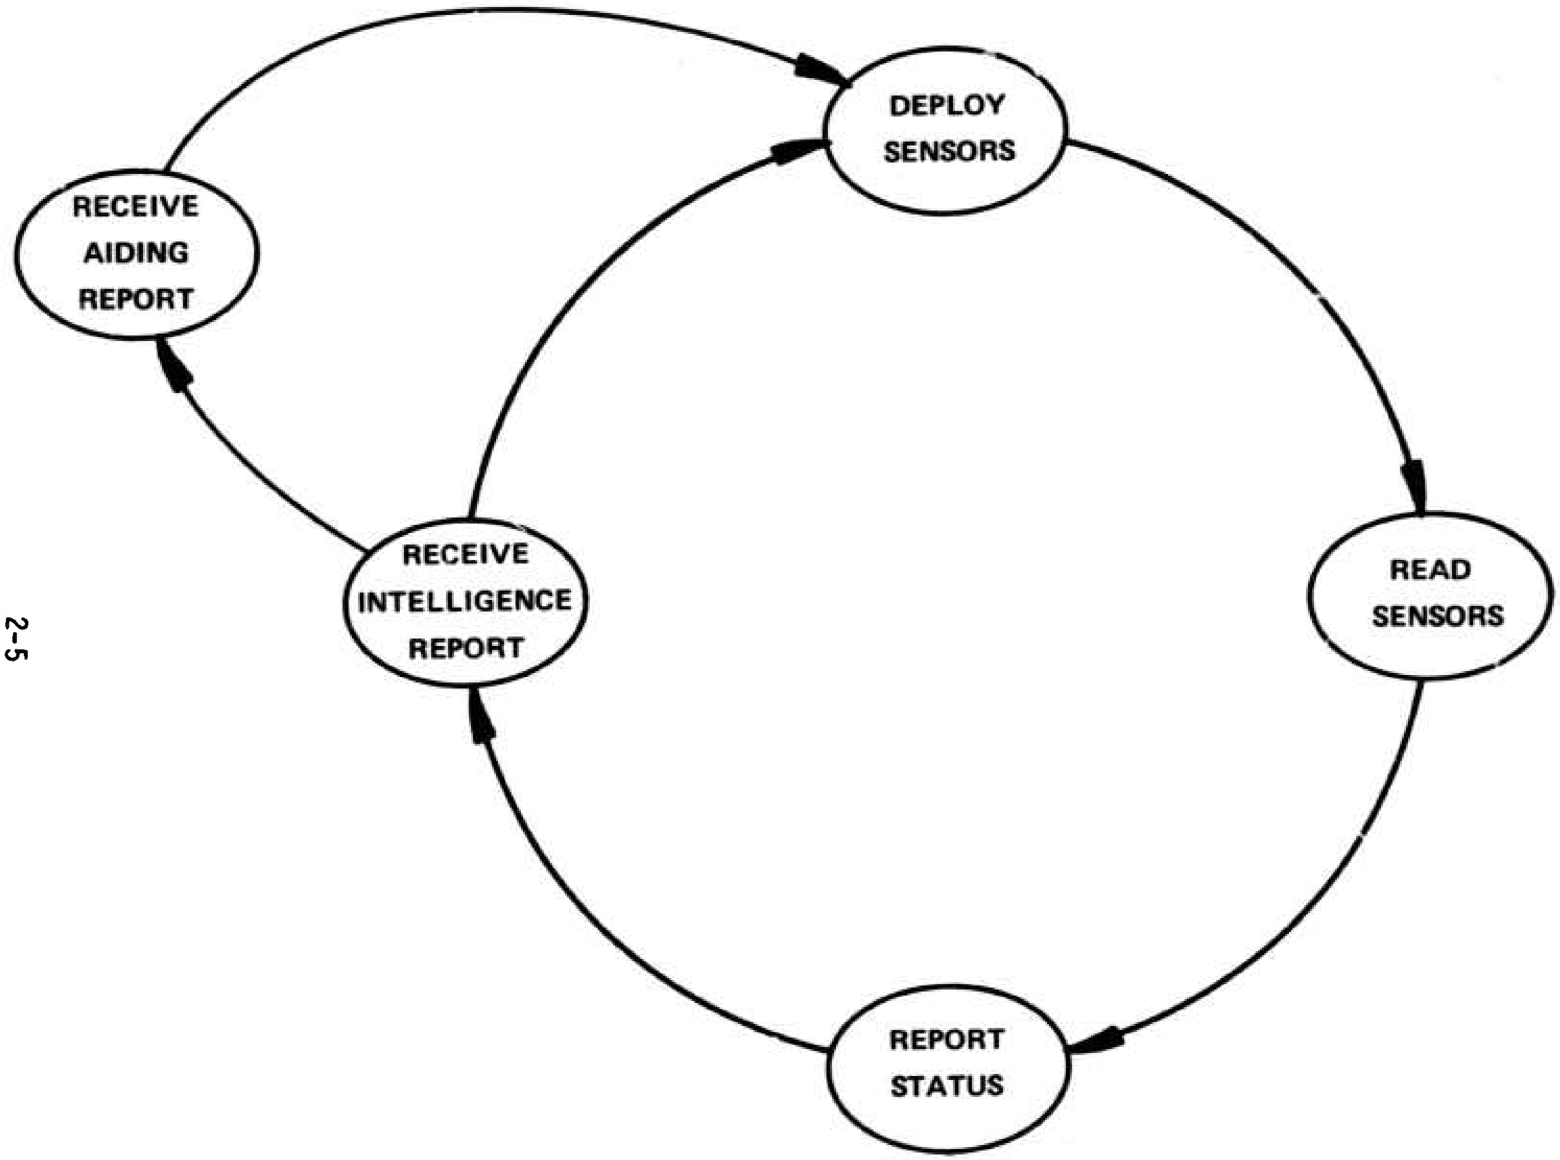

FIGURE 2-1. DECISION TASK SEQUENCE

Best Available

Copy

for all Pictures

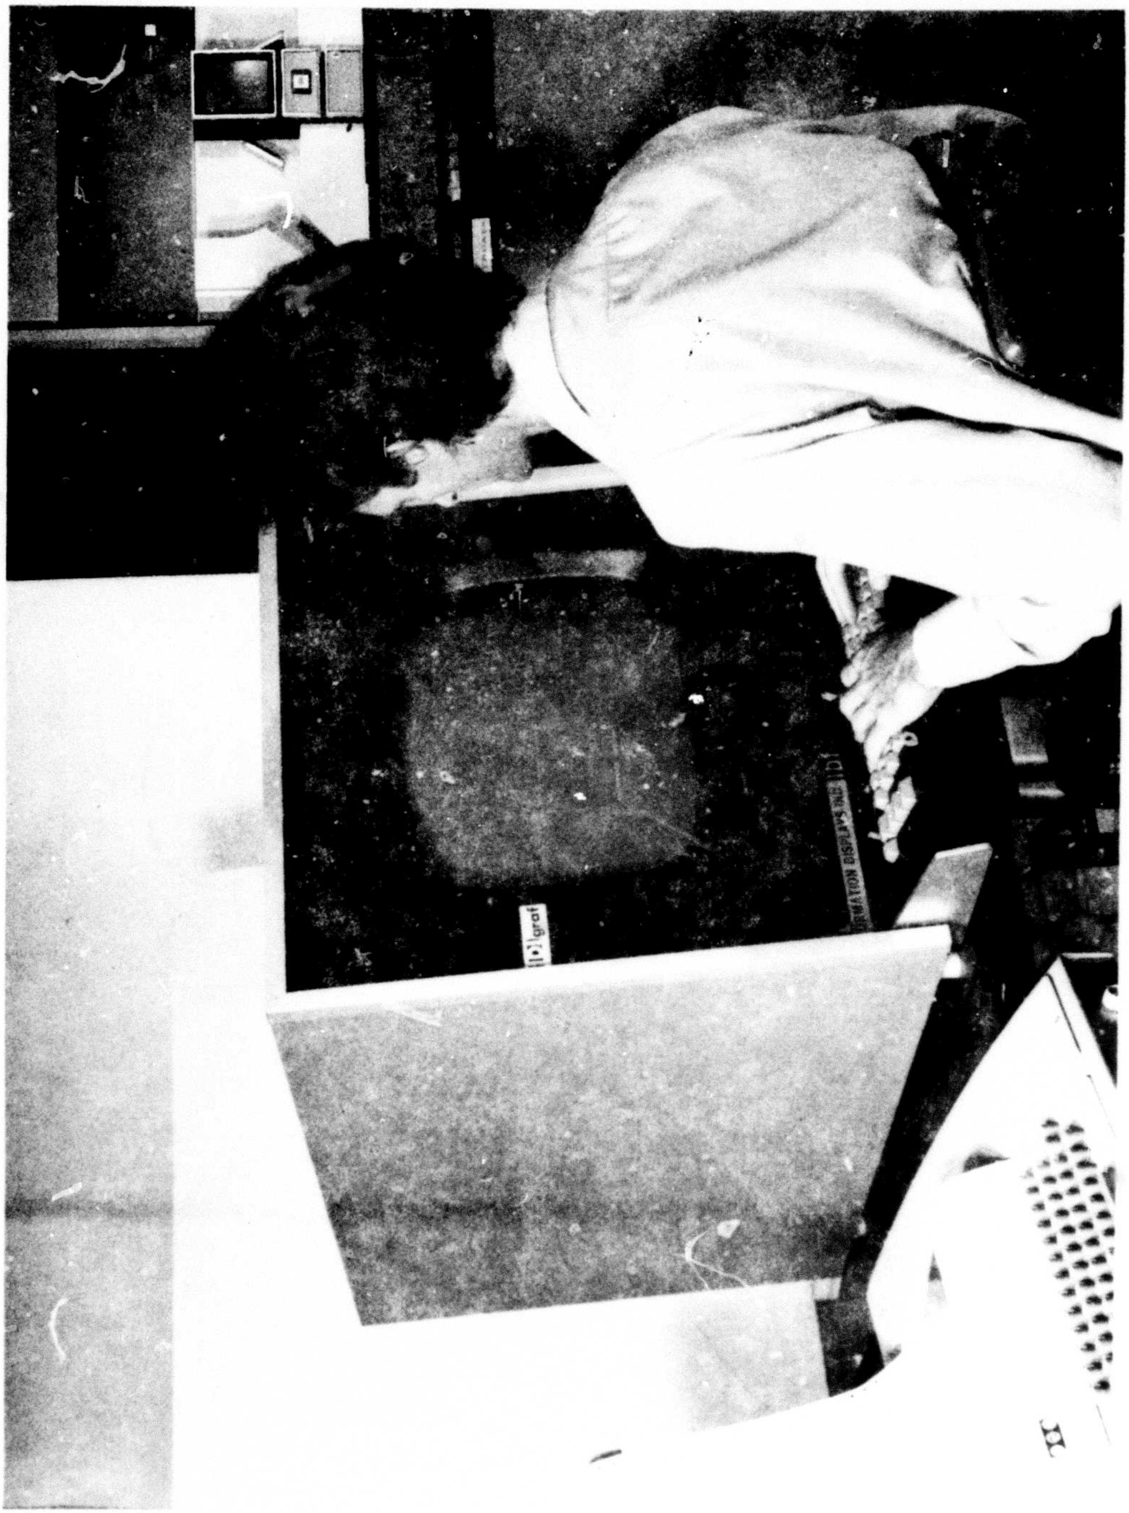

FIGURE 2-2. ADDAM SYSTEM TASK WORK STATION

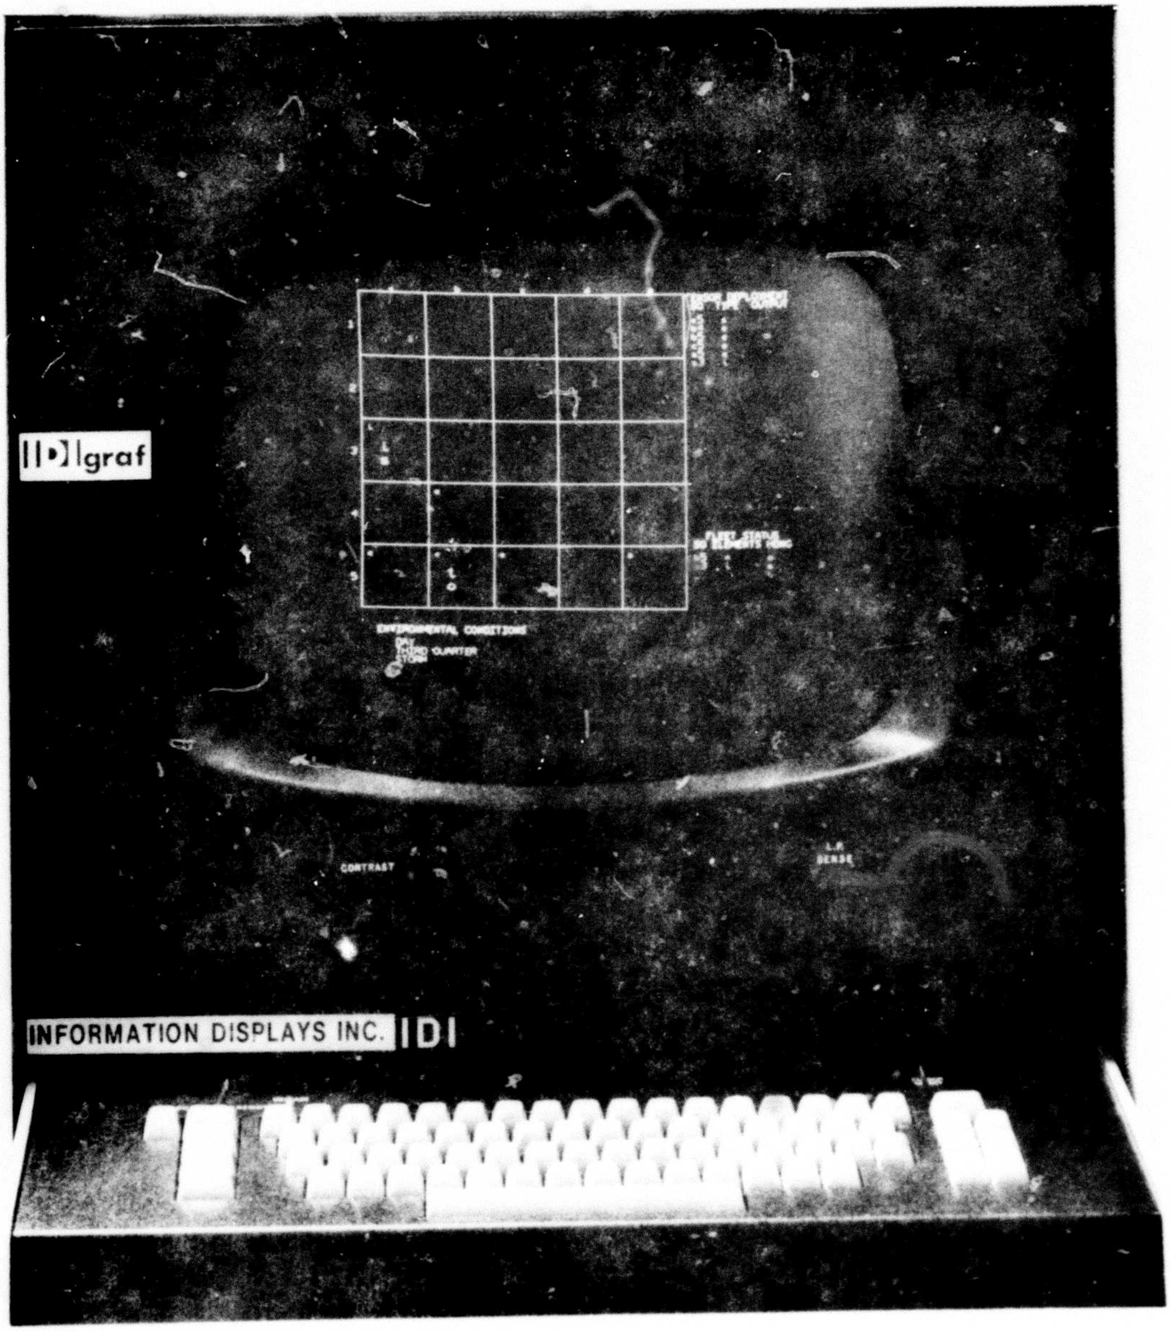

FIGURE 2-3. OPERATOR INFORMATION DISPLAY

location on the board. The report is generated in the abbreviated format shown in Figure 2-4 in order to speed up the printing on the teletype. Only those locations with non-zero probabilities appear in the report.

The intelligence analysis report is based on the information in the operator's status report and the current environmental conditions. The probabilities in the report are those which would be used to generate the next state of the environment if the operator's status report accurately represented the true state of the environment. Thus, the accuracy of the intelligence report depends on the accuracy of the operator's status report.

Decision Aiding. Decision aiding for the present experiment consists of recommendations to deploy sensors at various locations on the board. The recommendations are obtained by means of an adaptive maximum expected utility model of the decision process (see below). The model makes use of the dynamic estimates of the operator's utility for information from each type of sensor, the probabilities obtained from the intelligence analysis report, and the reliability and cost of each sensor.

The aiding information consists of a board location and a sensor type for each location on the board, but only those locations where the sensor type is non-null are displayed. The recommendations are displayed on the graphics display terminal. The operator can then go down the list, accepting or rejecting each recommendation. He can also supplement the list with additional sensor deployments.

2.2 The ADDAM System

Functional Organization. The functional organization of the ADDAM system is illustrated in Figure 2-5. The environment generator probabilistically generates the dynamic decision environment on the basis of expert probabilities and an organization structure specified by the experimenter.

INTELLIGENCE REPORT

	PROBABILITIES			
SQ	I	T	N	OBJ
A1	48	0	0	48
B1	25	0	0	25
E1	0	12	0	12
A2	25	0	0	25
D2	0	18	0	18
E2	0	1	49	50
E3	0	18	0	18

FIGURE 2-4. INTELLIGENCE ANALYSIS REPORT

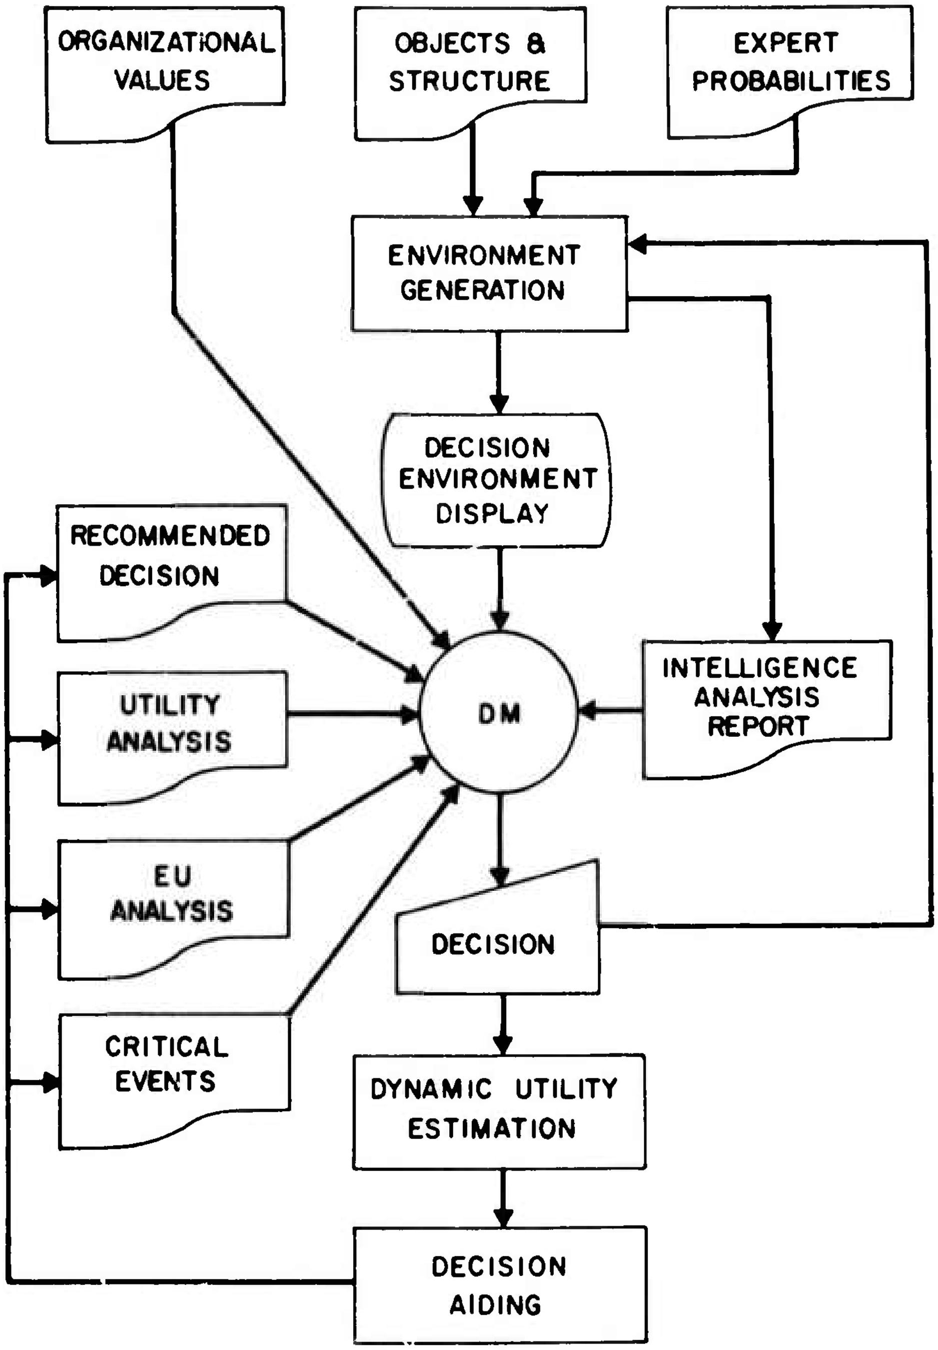

FIGURE 2-5. ADDAM FUNCTIONAL ORGANIZATION

The environment, as seen through sensors deployed by the operator, is displayed on a graphic display terminal. The operator makes decisions to deploy new sensors and to report on the status of the environment. These decisions are made on the basis of sensor information, the intelligence analysis report, organizational values (sensor costs, strategy instructions, etc.), and varying forms of decision aiding.

The operator's decision behavior is analyzed in order to dynamically estimate his utilities for intelligence information from the sensors. These utilities, estimated by using pattern classification techniques, are the basis for decision aiding. In the current study, the only form of aiding is recommended sensor decisions.

Adaptive Decision Model. The adaptive decision model is a model of the operator's decision behavior in deploying sensors to track objects in the environment (Freedy, Weisbrod, Davis, May, Weltman, 1974). A maximum expected utility model, based on the operator's utilities for information from each kind of sensor, is used. The utilities are adaptively estimated by using pattern classification techniques to track the operator's decision behavior.

The expected utility of deploying a sensor of type k at location L is the sum of the utilities of true positive and true negative sensor responses, minus the utilities of false positive and false negative responses and the cost of deploying the sensor:

$$EU_k(L) = \sum_i Q(p_i)\, M_{ik}[p_i(L) \cdot {}_kU_i^+ - (1-p_i(L))\, p_{\beta k}\, {}_kU_i^+ + (1-p_i(L))\,(1-p_{\beta k})\, {}_kU_i^-] - C_k \qquad (2\text{-}1)$$

where

M_{ik} = 1 if $i\epsilon\{i$: sensor k can report the presence of objects of type $i\}$

= 0 otherwise

$Q(p_i)$ = 0 if $p_i(L) = 0$

= 1 otherwise

$p_i(L)$ = p (object of type i at location L)

$p_{\beta k}$ = p (false positive from a sensor of type k)

${}_kU_i^+$ = Utility of a positive response for an object of type i by a type k sensor

${}_kU_i^-$ = Utility of a negative response for an object of type i by a type k sensor

C_k = Cost of deploying type k sensor

The utility matrix (Figure 2-6) is divided into two parts. One part contains the utilities for information that an object is present and the other contains the utilities for information that an object is not present. Since it is not possible to obtain information about trawlers or nets from an iceberg sensor, for example, the utilities for that kind of information are not represented.

It is impossible for the DM to distinguish between true and false alarms without additional information (and an additional decision). For this reason, the model only considers the actual sensor responses. However, the reliability of the sensor will affects its usage by the DM, and this will be reflected in the estimates of his utilities for information from that sensor.

OBJECT TYPE

$$
\begin{array}{cc}
 & \begin{array}{ccc} \text{ICEBERG} & \text{TRAWLER} & \text{NET} \end{array} \\
\text{SENSOR TYPE} \begin{array}{c} T1 \\ T2 \\ N \\ I \\ S \\ E \\ \hline T1 \\ T2 \\ N \\ I \\ S \\ E \end{array} &
\left[\begin{array}{ccc}
-- & {}_{T1}U^{+}_{T} & -- \\
-- & {}_{T2}U^{+}_{T} & -- \\
-- & -- & {}_{N}U^{+}_{N} \\
{}_{I}U^{+} & -- & -- \\
{}_{S}U^{+}_{I} & {}_{S}U^{+}_{T} & {}_{S}U^{+}_{N} \\
{}_{E}U^{+}_{I} & {}_{E}U^{+}_{T} & {}_{E}U^{+}_{N} \\
\hline
-- & {}_{T1}U^{-}_{T} & -- \\
-- & {}_{T2}U^{-}_{T} & -- \\
-- & -- & {}_{N}U^{-}_{N} \\
{}_{I}U^{-}_{I} & -- & -- \\
{}_{S}U_{I} & {}_{S}U^{-}_{T} & {}_{S}U^{-}_{N} \\
{}_{E}U^{-}_{I} & {}_{E}U^{-}_{T} & {}_{E}U^{-}_{N}
\end{array} \right]
\end{array}
$$

U^{+} MATRIX

Utility for Information that an Object is Present

U^{-} MATRIX

Utility for Information that an Object is Not Present

FIGURE 2-6. UTILITY MATRIX

Dynamic Utility Estimator. The dynamic utility estimator, schematically represented in Figure 2-7, is based on the principle of a multi-category pattern classifier whose discriminant functions are the expected utilities of each sensor. The pattern vector is

$$\overline{P} = [{}_1p_1, {}_1p_2, \ldots {}_kp_i \ldots]. \qquad (2\text{-}2)$$

Its components ${}_kp_i$ are a function of the probability that an object (fishing trawler, etc.) of type i is present and the reliability of the sensor, k. The pattern weights, which characterize the discriminant functions, correspond to the operator's utilities for information.

The utility estimator computes the EU of each sensor at each location on the board and selects those sensors (including a "null" sensor) for which the EU is maximum. The selected sensors are compared with the actual decisions made by the operator and if they differ the appropriate utilities are rewarded (increased) or punished (decreased) by the training procedure. Thus the utilities are trained to characterize the operator's judgmental behavior -- i.e., to make the utility estimator respond with the same decisions as the operator.

System Implementation. The ADDAM system is implemented on an Interdata Model 70 minicomputer with 24k bytes of core memory. An Information Displays, Inc. IDIgraf graphics display terminal is the primary interface with the operator. A teletype is used to provide printed output. A more complete description of the system can be found in Freedy, Weisbrod, Davis, May, and Weltman (1974).

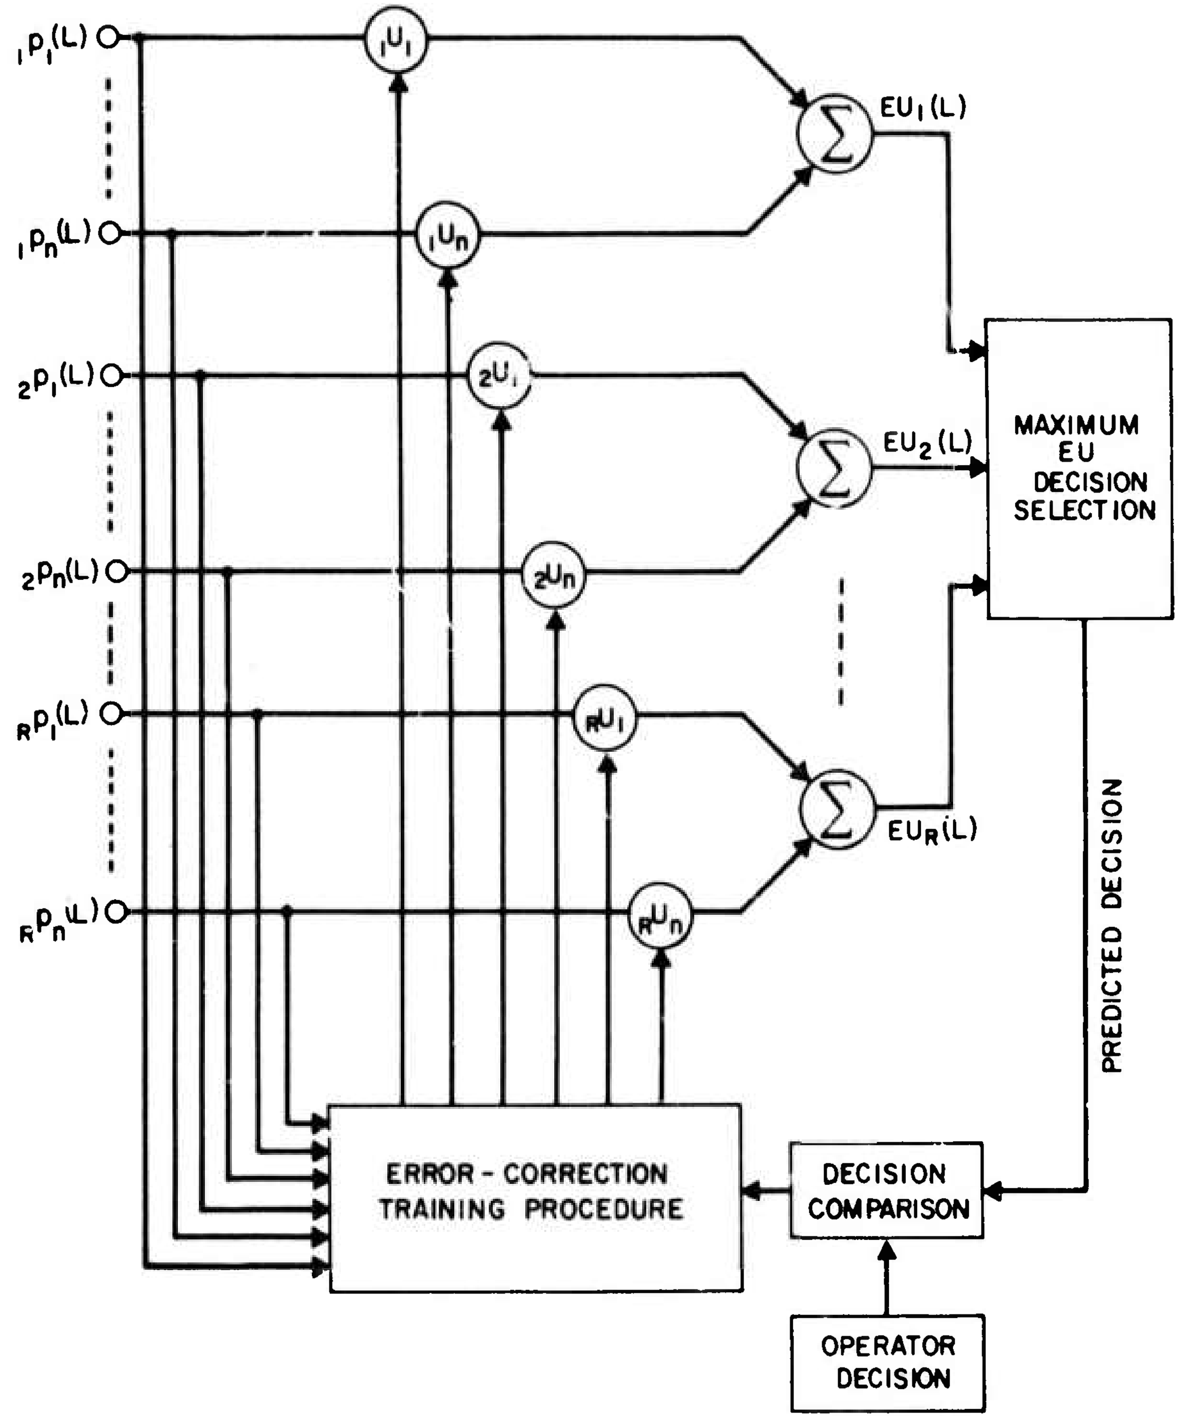

FIGURE 2-7. SCHEMATIC REPRESENTATION OF DYNAMIC UTILITY ESTIMATOR

3. PILOT STUDY

3.1 Objectives

The primary objective of the pilot study was to gain experience with the ADDAM System as a prelude to more formal experimental studies. In-house personnel gained considerable experience with the system and insights into its operational characteristics while performing shakedown tests and evaluations, but this experience did not include the reactions of naive operators (subjects) in an experimental setting. Specifically, the purpose of the pilot study was to obtain preliminary data on the following topics:

(a) Adequacy of instructions to the operator

(b) Operator response to task objectives and system interface

(c) Convergence of utilities under varying operator strategies and instructions

(d) Improvements to task and scenario programs.

3.2 Description

Subjects and Indoctrination. Three males between the age of 22-28 served as subjects. One subject failed to complete the experiment due to unexpected employment commitments. The subjects were instructed in the nature of the task. They were told to deploy sensors in order to locate and report on two objects, an iceberg and a trawler (which sometimes deployed nets). The intelligence report was described as a summary of the opinions of experts who, if told where the objects actually were, could predict where the objects would be during the next cycle.

The subjects were given a one-hour practice session in which they became familiar with the hardware and the input formats and abbreviations. The subjects learned rapidly. By the end of the first experimental session they were responding at their steady state decision rates. Following the last experimental session the subjects were given a debriefing questionnaire which sought to obtain information about the adequacy of the instructions, their reaction to the task, and any improvements they could suggest.

Strategies. The subjects were instructed to follow one of three strategy rules in performing the task. Strategy rules SR1 and SR2, summarized in Tables 3-1 and 3-2 respectively, represent organizational values externally imposed upon the operators. These rules specify the circumstances under which each type of sensor was to be used. Under SR1, if the probability of an iceberg (as reported by the intelligence report) is less than 20 percent, then an "*iceberg*" sensor is to be deployed, and if it is greater than 2 percent a "*something*" sensor is to be used. Similar rules define what to do on the basis of the probability of a trawler or a net. If there is a conflict between two or more rules, an *everything* (E) sensor is used. The SR1 rules are applied to each square for which the intelligence report probability is greater than percent. The strategy rules for SR2 are similar.

The third strategy, SR3, is a free strategy. No constraints on the use of sensors are imposed on the operator and he is free to define his own rules of behavior.

Experimental Plan. The experimental plan for the pilot study is summarized in Table 3-3. Each of the subjects were to be run under two sets of conditions. Subject 1 was run first under SR1 and then under FREE conditions.

TABLE 3-1

Strategy Rules SR1

PROBABILITY	FOR OBJECT I	PROBABILITY	FOR OBJECT T	FOR OBJECT N
2-20	Use i	2-50	Use e	Use n
21-100	Use s	51-70	Use t_1	Use n
		71-100	Use t_2	Use n

Use e sensor if there is a conflict

TABLE 3-2

Strategy Rules SR2

PROBABILITY	FOR OBJECT I	PROBABILITY	FOR OBJECT T	FOR OBJECT N
2-50	Use e	2-30	Use s	Use s
51-100	Use i	31-100	Use t_1	Use n

Use e sensor if there is a conflict

TABLE 3-3

Experimental Plan

SUBJECT	FIRST CONDITION	SECOND CONDITION
1	SR1	FREE
2	SR1	SR2
3	FREE	SR1

For subject three, the conditions were reversed. Subject 2 was run under SR1, but then dropped out of the experiment before he could run with SR2. Thus, SR2 was not actually used in the pilot study.

3.3 Results and Discussion

Firgures 3-1, 3-2, and 3-3 show how the utility estimates typically vary with time. Figure 3-1 shows one subject's utilities for information from the two kinds of trawler sensors and Figures 3-2 and 3-3 show the utilities for information from *something* and *everything* sensors, respectively. At time zero the utility estimator is untrained and assumes that all utilities are equal. An initial value of 1.00 was arbitrarily chosen. As the training progresses, the utility estimates separate into distinct levels and tend to stabilize at these levels (i.e., they converge). The figures are plotted with straight line connections between the data points. This was done to facilitate the visual separation of the individual curves. Actually, the utility estimates increase or decrease in a stepwise, rather than continuous, fashion.

The utilities for information from *something* and *everything* sensors (Figure 3-1) clearly show their usage preferences. According to strategy S1, *something* sensors are used to track icebergs, but never trawlers or nets. The utilities for information about trawlers and nets are virtually in-distinguishable from each other, but clearly separated from the utility for

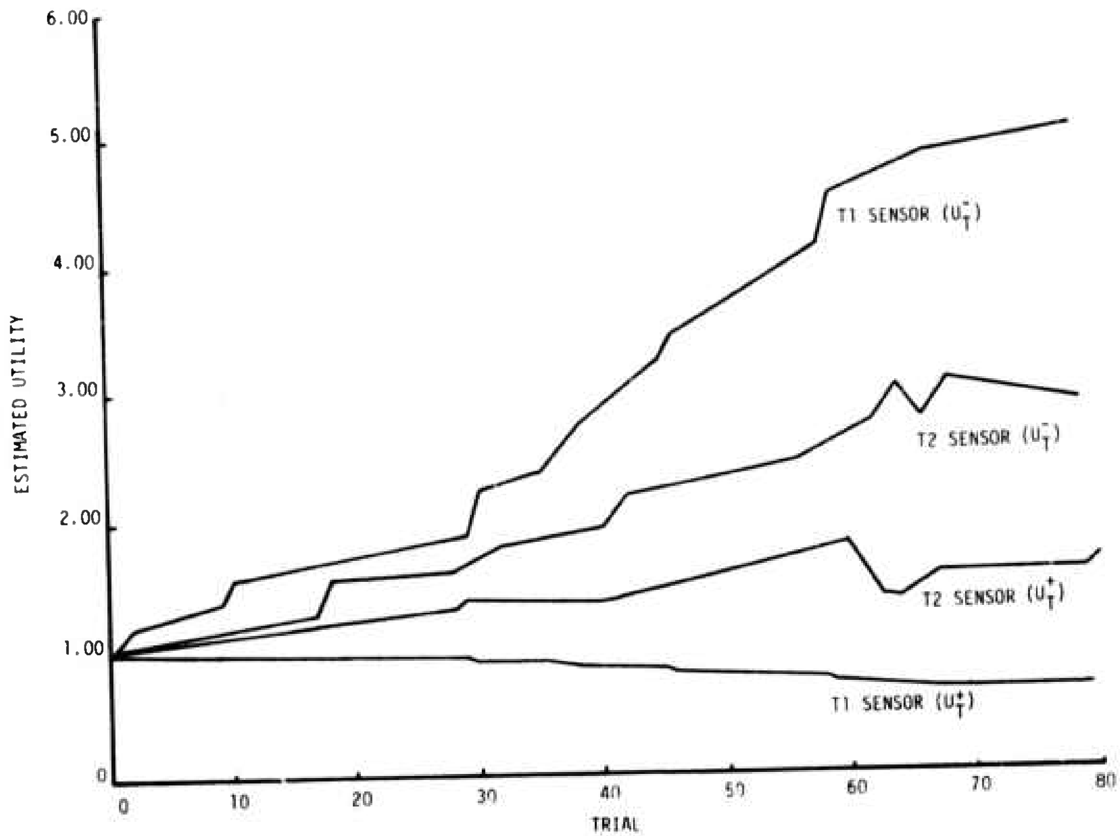

FIGURE 3-1. ESTIMATED UTILITIES FOR INFORMATION FROM TRAWLER SENSORS

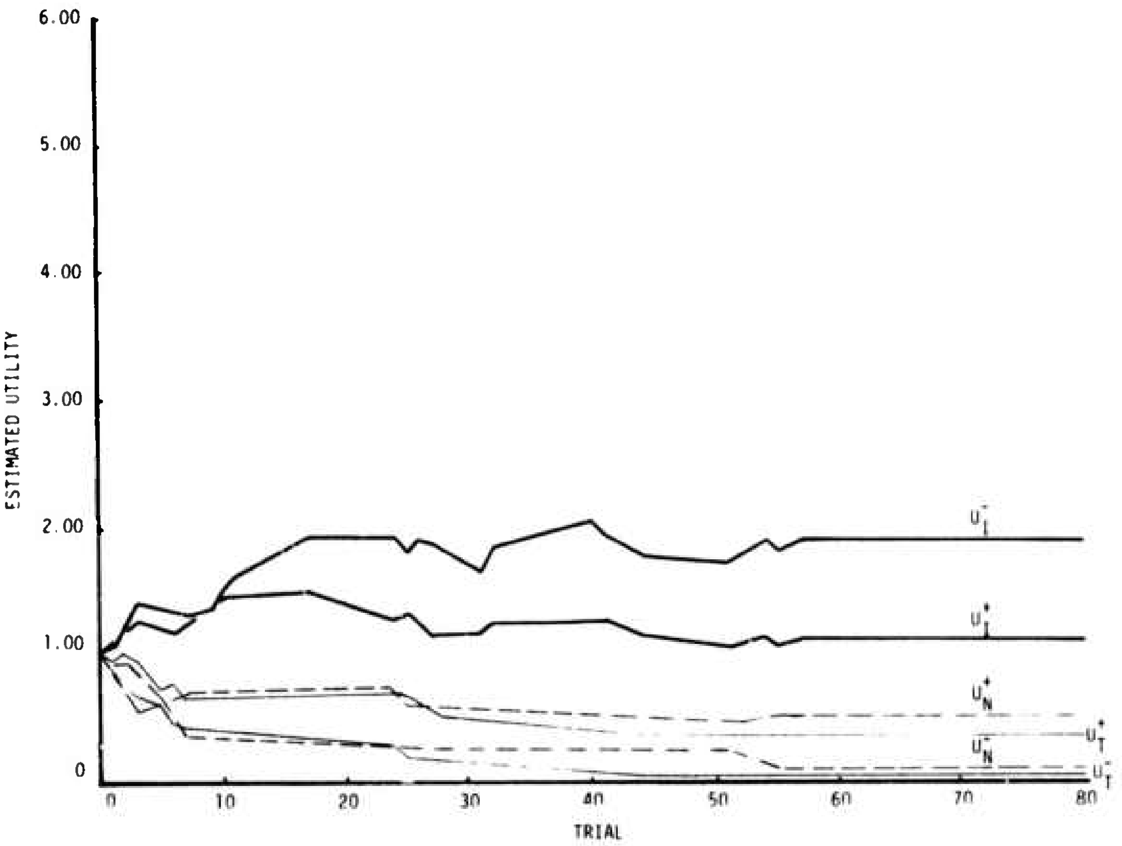

FIGURE 3-2. ESTIMATED UTILITIES FOR INFORMATION FROM SOMETHING SENSORS

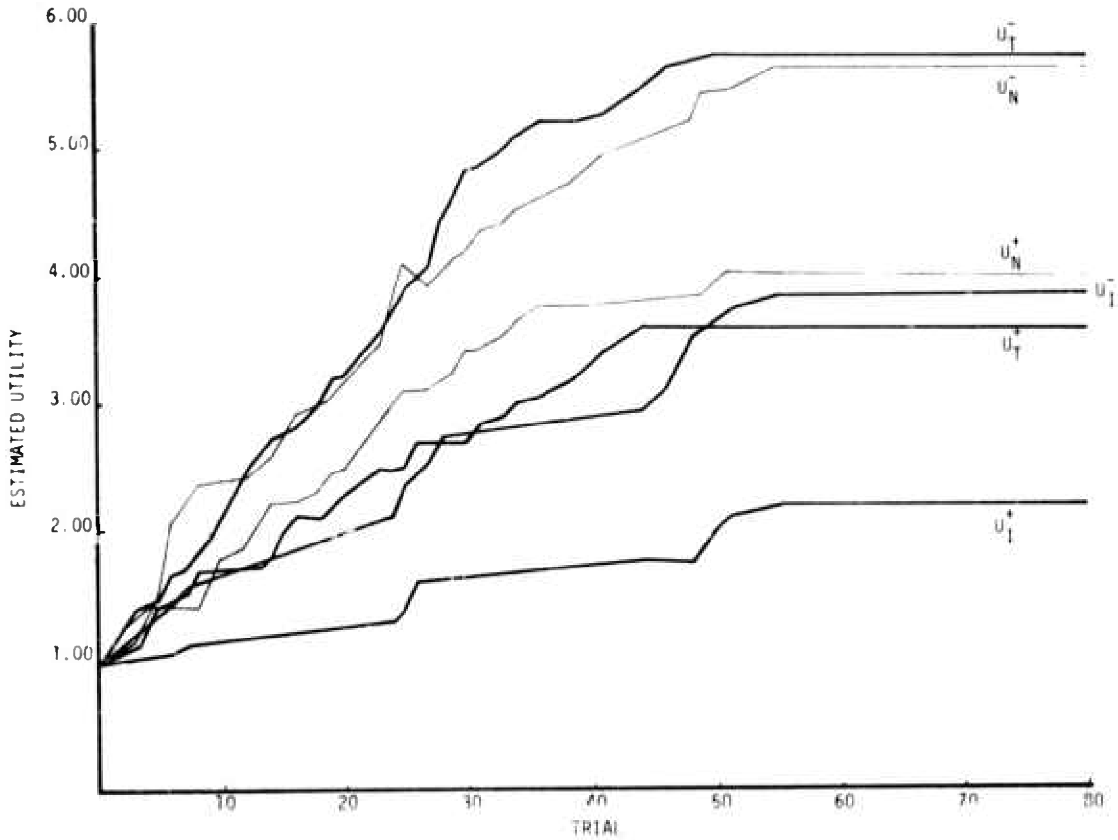

FIGURE 3-3. ESTIMATED UTILITIES FOR INFORMATION FROM EVERYTHING SENSORS

information about icebergs. Similarly, *everything* sensors are used for detecting trawlers and, to a lesser extent nets (wherever there is a probability of both a trawler and a net). They are infrequently used when icebergs are present (only when other objects have a probability of occurring). This pattern is reflected in the utilities for information from an *everything* sensor.

During the pilot study, ADDAM proved to be very erratic in predicting the operator's decisions. A particular sensor might be used a few times, especially in low probability situations, causing small incremental jumps in the values of its U^+ utilities and large jumps in its U^- utilities. This would result in a prediction that the sensor will be deployed at most locations on the board, especially those where the probability of finding an object was zero. Subsequently, the U^- utilities would be heavily punished and no predictions would be made, even though the operator used the sensor under certain circumstances. The utilities would then be rewarded until the cycle repeated itself again several trials later. Because of this erratic behavior and because of the strong interactions between the utility values for different sensors which characterize the adaptive training mechanism, no comparisons of utilities for obtaining the same information from different sensors were made.

The erratic behavior of the U^- utilities was caused by a singularity in the adaptive EU model which occurred when the probability of an object was zero. This singularity was corrected by introducing the Q-term (see equation 2-1) to the model. This term essentially switches off parts of the model when the probability is zero. During subsequent shakedown tests, the corrected model was able to predict the SR1 strategy approximately 95% of the time once it was trained.

The subjects were given debriefing questionnaires following their last experimental session. They all seemed to agree that the instructions were adequate and that they understood the task by the end of the first experimental session. The subjects found the task very interesting and became quite involved in performing it.

Both written and verbal comments were made concerning the subject's lack of tolerance of ambiguities in the task situation. These ambiguities resulted from the requirement that the subject report status after every trial. During the experiment a subject would sometimes find that he had been tracking a series of false alarms and, as a result, had completely lost track of the object. Because he was required to make a status report, the subject would make a wild guess about the location of the object. These guesses would only occasionally be confirmed, usually by additional false alarms. Since the intelligence report was also based on the incorrectly reported status, the resulting degree of ambiguity was apparently at the threshold of tolerance.

The problem was alleviated during the full scale experiment by simply not requiring a status report after each trial. If a subject loses an object he is tracking, he does not make a status report for that object and, consequently, does not receive an intelligence report for that object. He then begins to search for the object. The utility estimation algorithm is able to recognize that this situation is different from normal tracking behavior and does not train the subject's utilities while he is in "search mode".

The pilot study helped identify a number of necessary changes and improvements to the ADDAM system. Some of the changes were implemented immediately for use in the full scale experiment while others are currently being implemented and will affect future experiments. The immediate changes included modifications to the model and training algorithm to deal with the singularity at zero probabilities and related changes in the intelligence report generator. Long term changes include the addition of facilities for automatic recording and reporting of internal system processes for use as experimental data or computation and decision aiding display of cost and payoff information to subjects (at experimenter option), display of real world feedback for sensor validation (at experimenter option), and automatic computation of convergence measures. Other changes are the addition of new sensors for a more balanced game and modifications to improve the sensitivity of the utility training algorithm.

4. INITIAL STUDY OF CONVERGENCE AND DECISION AIDING

4.1 Objectives

The overall goals of this research program are: (1) to implement promising techniques for adaptive modeling of human decision behavior; (2) to explore in experimentally controlled environments the factors which influence the effective monitoring, aiding, and automating of dynamic decision processes; and (3) to establish guidelines for the application of adaptive decision systems. The purpose of this initial study was to provide preliminary data on the ability of the system to adaptively acquire decision strategies, to predict operator behavior, and to aid the operator in making decisions.

The experiment had three major objectives. The first was to validate the utility estimates and the adaptive decision model. This was done by demonstrating that the model is capable of predicting decision maker (DM) behavior with a reasonable degree of accuracy in the absence of any aiding to the DM. It was also demonstrated that the model was sensitive to individual DM decision strategies.

The second objective was to determine some of the major effects of aiding on the consistency of decision making. This was done by examining the effect of decision aiding on the variability of the operator's utilities.

The third objective was to determine what effect the subject's apparent degree of control has on his acceptance of decision aiding. This was accomplished by investigating the effect of explaining to the subject his control over the kind of aiding he received.

4.2 Hypotheses

The hypotheses tested in the experiment were as follows:

(1) The adaptive expected utility model will predict decision maker behavior to a reasonable degree of accuracy in the absence of decision aiding to the subject.

(2) The model will be sensitive to individual variations in decision strategy.

(3) Aiding information in the form of sensor deployment recommendations will cause the utility estimates to converge within narrower limits.

(4) Aiding information in the form of sensor deployment recommendations will allow the subject to make more decisions per unit time.

(5) The degree to which subjects understand the adaptive aiding and their contribution to it (i.e., their perceived degree of apparent control) will result in more frequent acceptance of aiding and, concurrently, more stable utility estimates.

4.3 Experimental Design

A one-way experimental design with three treatment levels was used. This design is illustrated in Table 4-1. The treatments were based on levels of aiding and degrees of knowledge. The treatments were no aiding (control group), aiding without indoctrination, and aiding with indoctrination.

TABLE 4-1

EXPERIMENTAL DESIGN

GROUP	SESSION 1	2	3	4
I. NO AIDING	TRAINING	TRAINING	TRAINING	NO AIDING
II. AIDING WITH LOW APPARENT CONTROL	TRAINING	TRAINING	TRAINING	AIDING
III. AIDING WITH HIGH APPARENT CONTROL	TRAINING	TRAINING	TRAINING	INDOCTRINATION AND AIDING

The measures of interest in the experiment are the utility values, as calculated by the utility estimation algorithm; the utility matrix difference (UMD) score, a measure of the variability of the utilities; frequency of correct model predictions; and the number of decisions made by the subject per unit time.

4.4 Subjects and Procedures

Nine male subjects, three in each group, were recruited from two local colleges. Their ages ranged from 20 to 26 years and all were undergraduates. This sample resembles quite closely the potential users of computer-aided decision and control devices in the military.

Each subject had four sessions of 1-1/2 hours duration. The first three sessions were training sessions. The training included instructions on system operation, "hands on" experience with the equipment to familiarize the subjects with the input formats and other task features, and, in general terms, strategy instructions. The training procedure was the same for all subjects.

After the third session, Group III subjects were given an explanation of the adaptive nature of the decision model and how its sensor recommendations (aiding) were actually controlled by the subjects' own behavior. Group

II subjects did not receive such indoctrination and, thus, were not aware of the adaptive nature of the decision aiding they received. Group III subjects, thus, were presumed to have a high degree of apparent control over their aiding and Group II subjects a low degree. Group I subjects served as a control group. They received no indoctrination and no aiding.

The subjects in this experiment were paid on an hourly basis and were told they would receive a bonus based upon their performance.

4.5 Experimental Task

The task performed by the subjects was the same as described in Section 2.1 with one important change which resulted from the findings of the pilot study. The subject was no longer required to give a status report on every trial. If he lost track of an object he did not report its status and, consequently did not receive an intelligence analysis report on it.

The subjects were informed that they were participating in human decision making research. Their task was to accurately track the objects by placing sensors, evaluating the sensor reports, and declaring the location of the objects. They were told their performance would be measured in terms of the cost of accurately tracking the objects. Also, that their final performance would be compared with the other participants and they would receive pay bonuses depending on how well they did.

Once they had located and declared the status of the objects they would receive an intelligence report which would give the likelihood of the object(s) next move. They were told that the intelligence report was extremely accurate if they had accurately reported the object(s) location, i.e., the "expert" who generates the intelligence report has no knowledge of the location of the objects but knows with great accuracy how the objects behave.

The participants were instructed to deploy sensor according to the general strategy described in Table 4-2. This strategy refers to the probability values in the intelligence reports. Part of their decision task was to decide what probability values in the intelligence report would be considered high and low. Another aspect of their task, in the case of multiple reports, was to decide which of the sensor reports were false alarms.

TABLE 4-2

Generalized Strategy Rules

PROBABILITY OF ICEBERG	SENSOR TYPE	PROBABILITY OF TRAWLER	SENSOR TYPE	PROBABILITY OF NET	SENSOR TYPE
Low	i	Low	t1	Low	n
High	s	High	t2	High	n

Use the e sensor if there is the possibility that more than one object will be in the same location.

5. RESULTS AND DISCUSSION

5.1 Utility Measurement

Utility estimates are measures of subjective quantities which characterize a person's judgments, and they are valid only to the extent that they approximate these quantities (Peterson, 1971). Utility measures represent an abstraction of human judgment. As such, they have meaning only within the context of a model of human judgmental behavior. If the model accurately characterizes human decision making in a decision task, then within that context the utility estimates are intrinsically valid. With this in mind, two questions related to validity were examined during the experiment: (1) Does the adaptive expected utility model accurately predict decision behavior? (2) Are the dynamic utility estimates sensitive to individual variations in decision behavior?

The adaptive nature of the decision model used by ADDAM provides a means of examining the predictive validity of the utility estimates and of the model. This predictive validity is a matter of degree. That is, perfect predictive validity would require that the model's predictions be completely consistent with the operator's decision behavior (or vice versa). Under such circumstances, the utility estimates would each converge to a single value as the adaptive model learned the operator's behavior. If the operator behaves "most of the time" in a manner which is consistent with an expected utility (EU) model, the utility estimates will converge as the model learns, but there will be some variability in the steady state estimates. However, if the operator's behavior is inconsistent with an EU model, or highly "erratic", there will be a high degree of variability. For a task as complex as intelligence gathering, it is highly unlikely that a human operator would be perfectly consistent with an expected utility model.

Utility Training. Figure 5-1 shows a typical plot of utility estimates as a function of trial for the first three sessions of the experiment. The utilities are for information from an *iceberg* sensor. The I^+ curve is the utility for information that an iceberg is present and the I^- curve is the utility for information that no iceberg is present. During the first session, the subject (BH) is learning the task and the adaptive utility estimator is learning the subject's behavior, i.e., the adaptive mechanism is being trained. By the twentieth trial (during the second session) the utilities have begun to converge and only minor changes take place. These minor changes are due mainly to inconsistencies in the operator's behavior.

Predictive Accuracy. A comparison of the decisions of the control group subjects (those who received no aiding from the model) with the model's predictions during the fourth session indicate that the trained adaptive decision model is highly accurate in predicting subject behavior. Table 5-1 shows that the model predicted more than 95% of the sensors actually deployed by the subjects.

TABLE 5-1

Model Predictions of Sensors Actually Deployed by Control Subjects During the Fourth Session

SUBJECT	PERCENT OF SENSOR DEPLOYMENTS PREDICTED
AK	97
BH	97
GB	95

Utility Separation. The decision task is structured so that several different sensors could be used to track the same object. Because the sensors differ in reliability, cost, and usage pattern, it is to be expected that the

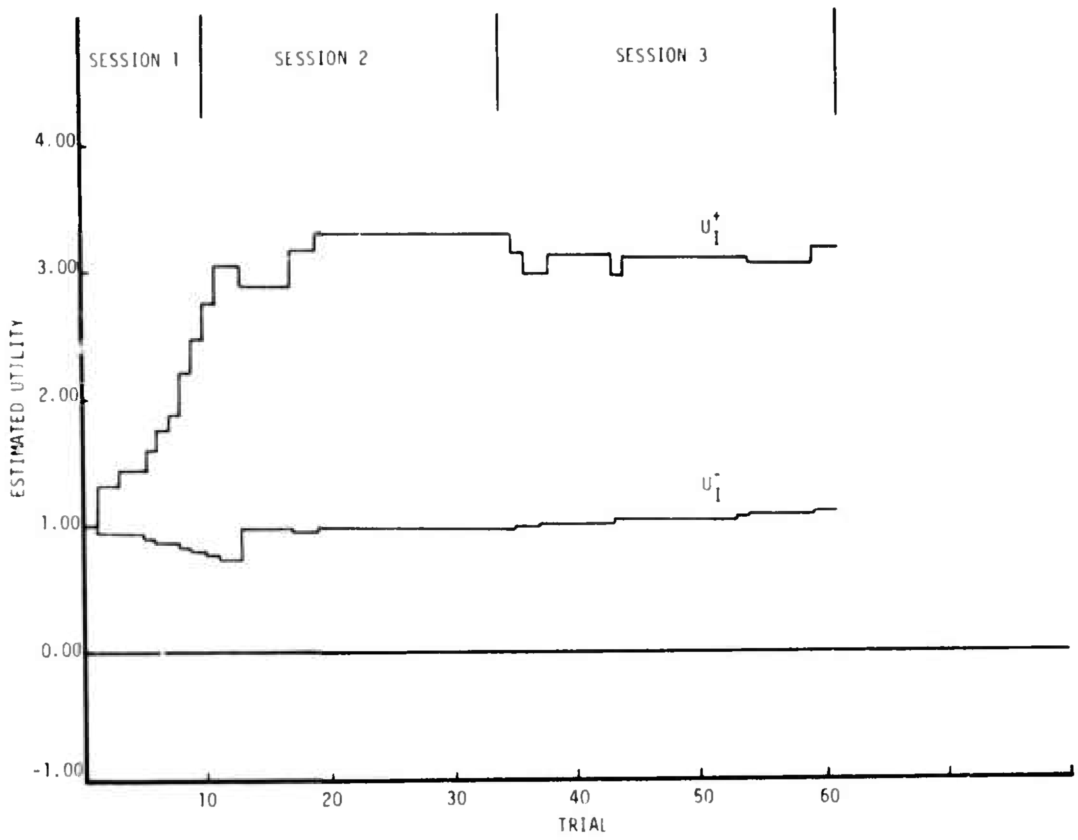

FIGURE 5-1. ESTIMATED UTILITIES FOR INFORMATION FROM ICEBERG SENSORS (SUBJECT BH)

estimated utility of the same information would differ from sensor to sensor. Figures 5-2 and 5-3 show one subject's utilities for information about icebergs from three different kinds of sensors. Figure 5-2 shows the estimates of the utilities of information that an iceberg was present and Figure 5-3 shows the utilities for information that an iceberg was not present. Only the values at the end of each session are plotted. These figures are typical of the separations achieved between the utilities for information from different kinds of sensors.

Indifference. One characteristic of the operator's behavior which is of particular interest is the point at which he is indifferent between two sensors which can give similar information. This indifference point corresponds to the intersection between the expected utility functions of the two sensors and can be computed from the utility estimates which have been adaptively derived from observations of the operator's decision behavior.

One indifference point is the intersection between the EU functions for the *iceberg* and *something* sensors. Figure 5-4 illustrates the expected utilities of these two sensors as a function of the probability of an iceberg. The EU functions are based on the utility estimates for subject JF at one point during the third training session. The probabilities of trawlers and nets are assumed to be zero. The point of indifference is at $p_I = 0.45$. When $p_I < 0.45$ the EU of an *iceberg* sensor is higher than that of a *something* sensor and the *iceberg* sensor is preferred. Similarly, when $p_I > 0.45$ the EU of a *something* sensor is higher and it is preferred.

The indifference point is determined by the value of the utility estimates at any given time. These values vary as the subject performs the task, therefore, the indifference points also vary. Figure 5-5 is a plot of the variations in the *iceberg/something* sensor indifference point during the third session for subject JF.

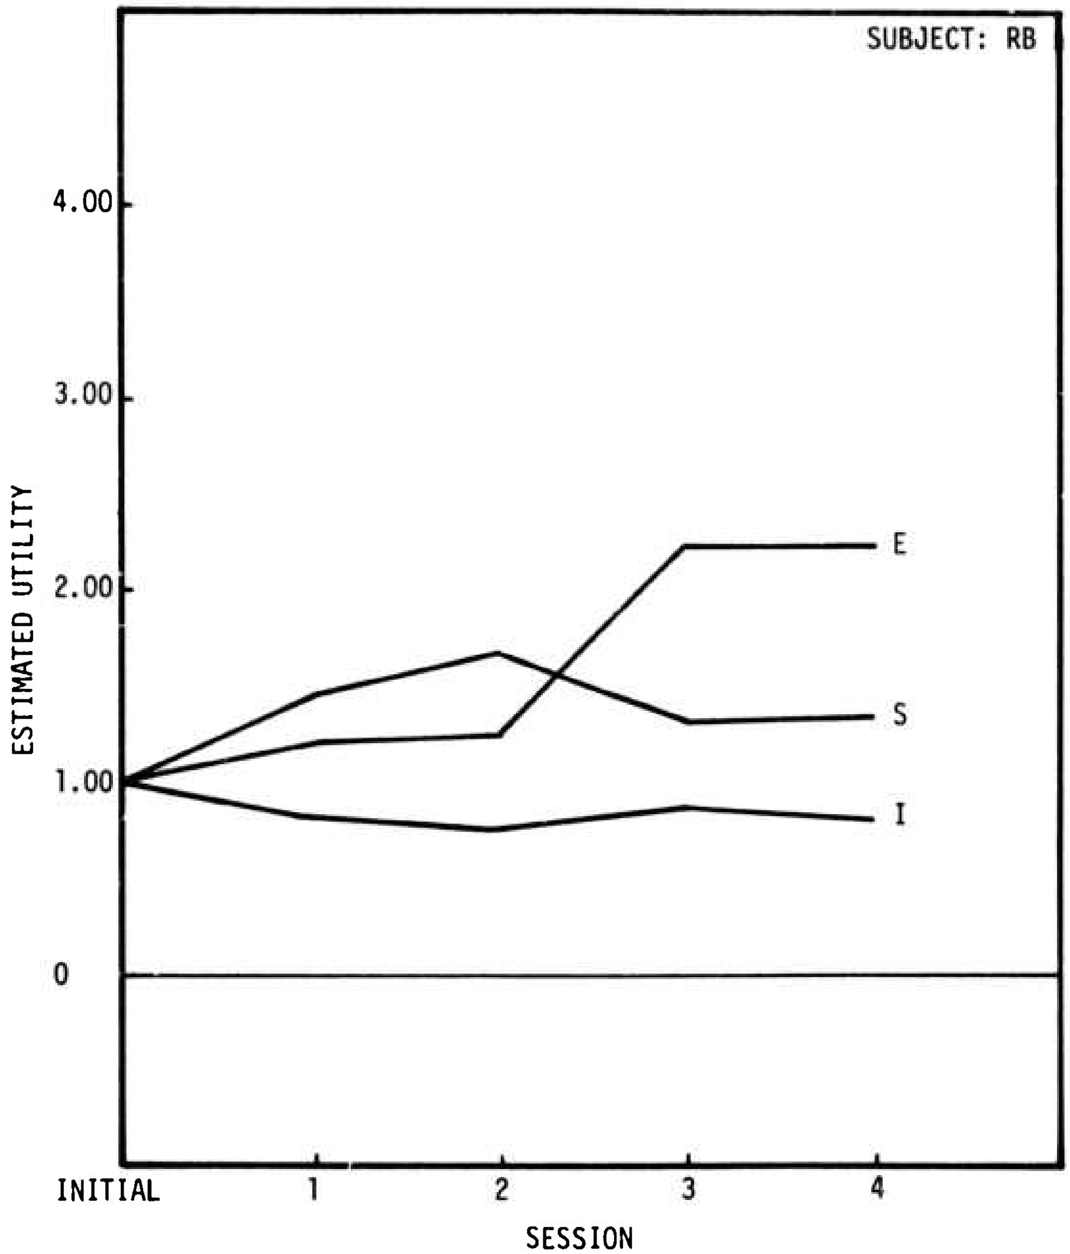

FIGURE 5-2. UTILITY FOR INFORMATION THAT AN ICEBERG IS PRESENT, FROM SEVERAL TYPES OF SENSORS

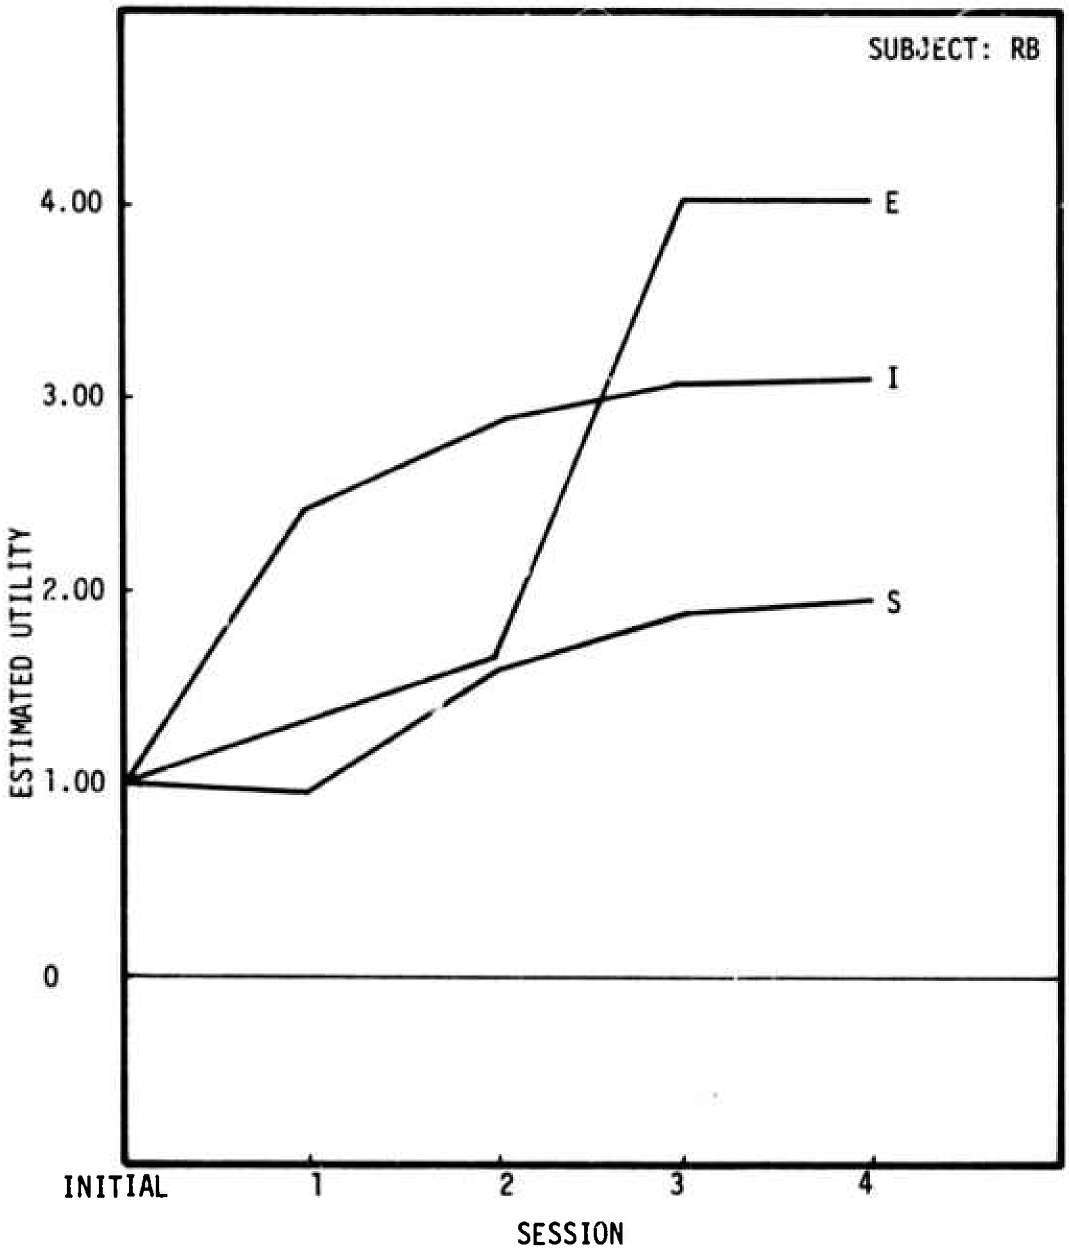

FIGURE 5-3. UTILITY FOR INFORMATION THAT AN ICEBERG IS NOT PRESENT, FROM SEVERAL TYPES OF SENSORS

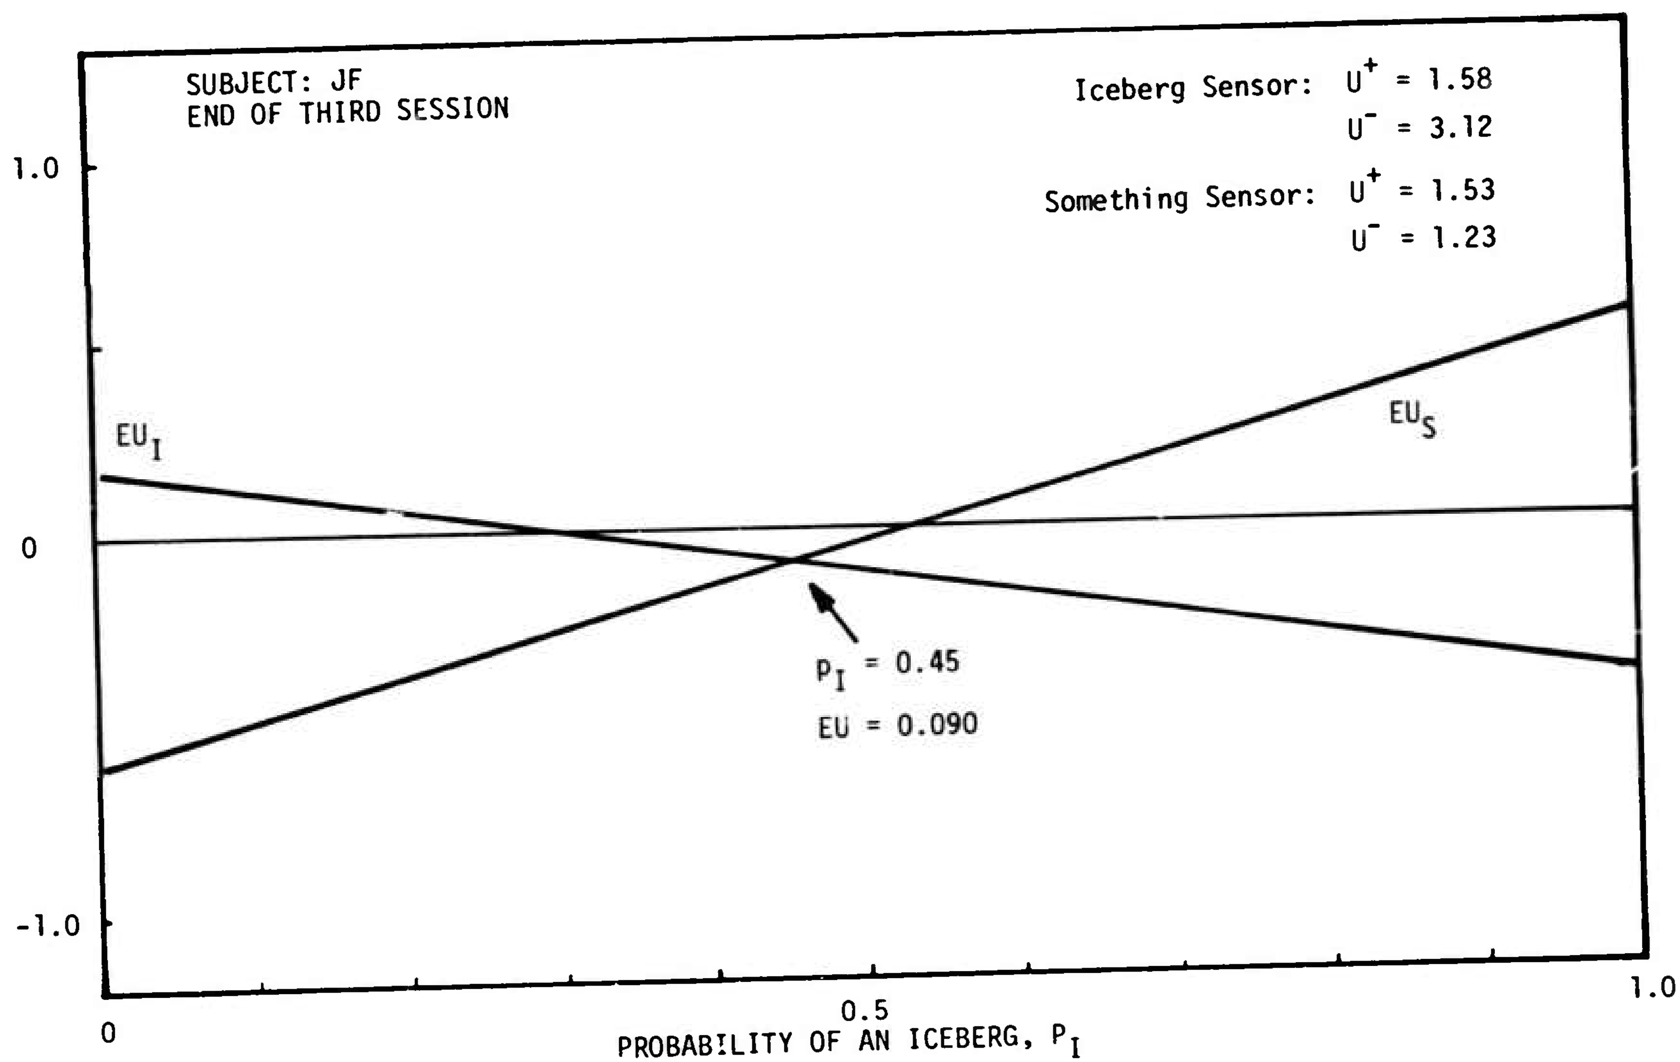

FIGURE 5-4. INDIFFERENCE POINT FOR ICEBERG AND SOMETHING SENSORS
($p_T = p_N = 0$)

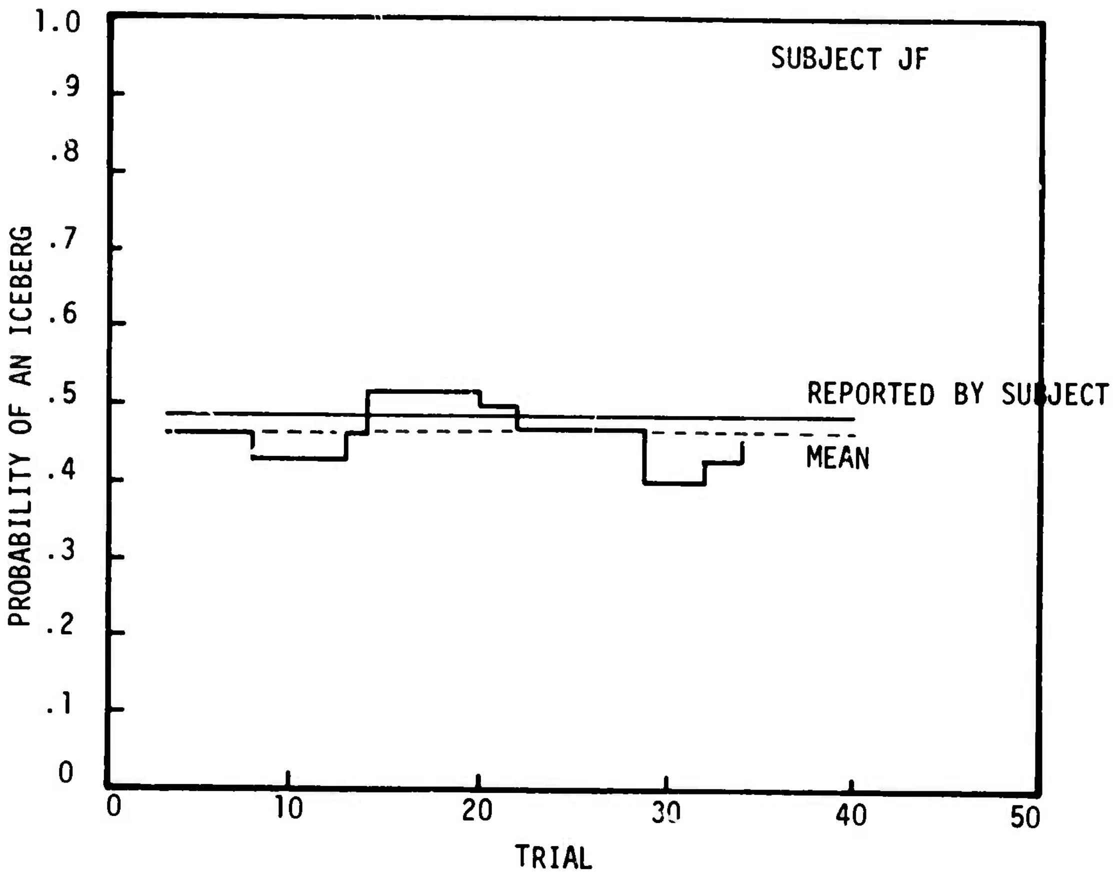

FIGURE 5-5. VARIATION IN INDIFFERENCE POINT FOR ICEBERG AND SOMETHING SENSORS BASED ON UTILITY ESTIMATES (FOR THIRD SESSION)

At the end of the third session, each subject was asked the following question: If you were forced to, what single probability number would you choose above which you would consider the probability of an iceberg to be high (and thus deploy a *something* sensor) and below which you would consider the probability to be low (and thus use an *iceberg* sensor). The reported *iceberg/something* sensor indifference point for each subject is tabulated in Table 5-2 as is the mean value of the indifference point computed on the basis of the third session utilities. These values are also plotted in Figure 5-6. The Pearson correlation between these two variables is 0.82, which is significant at the 0.01 level.

5.2 Decision Aiding

During the first three sessions of the experiment the utility estimates are adaptively adjusted so that the decision model conforms to the operator's decision behavior. The convergence of the utilities is an indicator of how well the model describes the operator's behavior. During the fourth session, however, the model is used to provide normative information (decision aiding) to the the operator and the meaning of convergence changes.

In evaluating the effects of decision aiding on the operator's decision behavior, we are essentially asking how well the operator is adhering to his decision model. During training, the operator's behavior is the standard to which the decision model is adjusted. During aiding, the decision model becomes a standard to which we compare the operator's behavior. During this period, variations in the utility estimates reflect the operator's reaction to the preferred aiding. If he accepts the aiding, that is, if he continued to act in accord with the machine's model of his previous decisions, no changes are made to the utilities. If he rejects the aiding, the utilitics are retrained.

TABLE 5-2

PREDICTED VS. STATED INDIFFERENCE POINT BETWEEN ICEBERG AND SOMETHING SENSORS

	SUBJECT	MEAN PREDICTED INDIFFERENCE POINT	STATED INDIFFERENCE POINT
	GB	.211	.30
G1	BH	.260	.30
	KK	.266	.25
	DK	.421	.50
G2	JL	.245	.30
	RB	.160	.20
	RF	.419	.30
G3	DW	.249	.30
	JF	.459	.48

Pearson Correlation .82

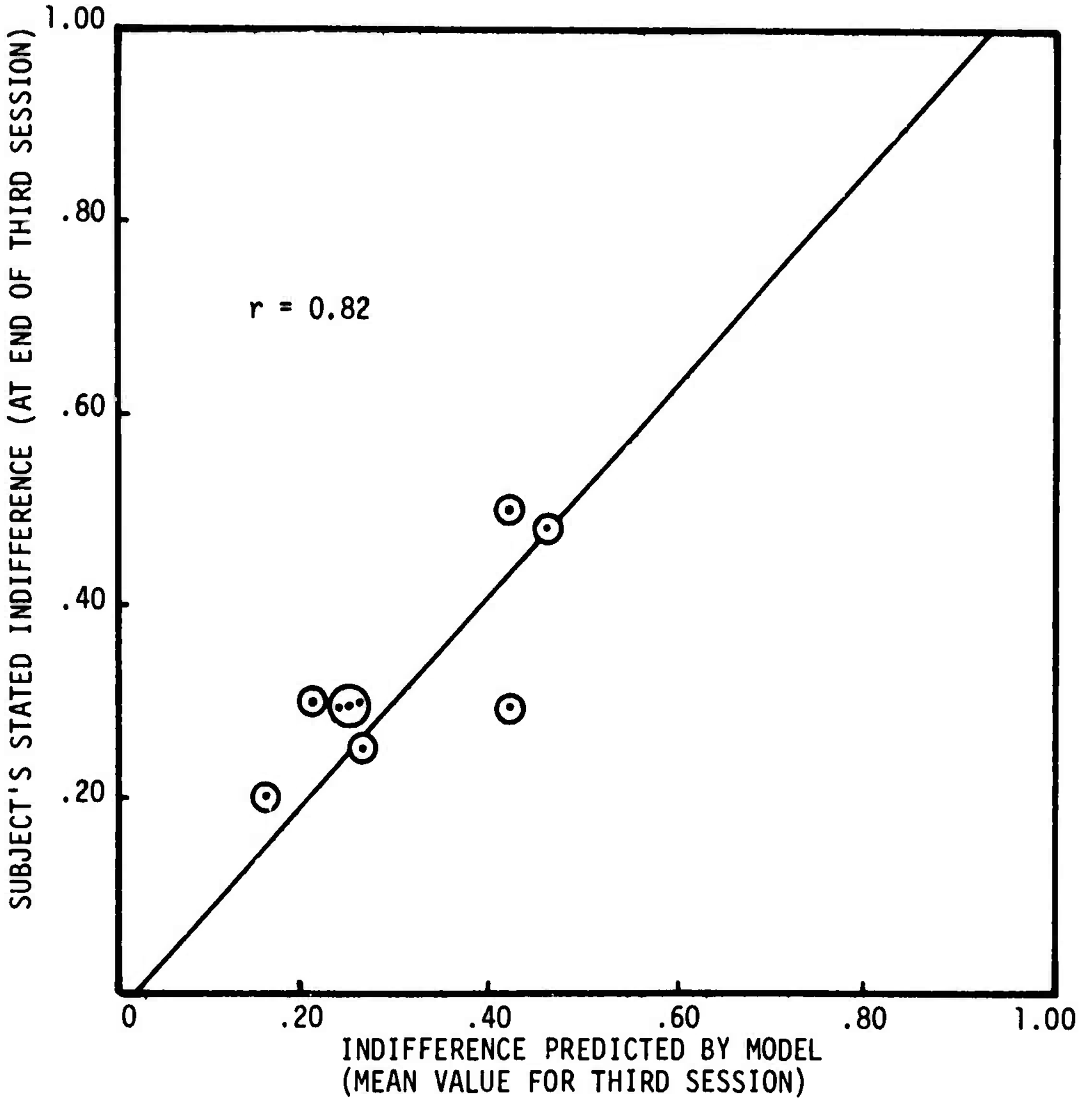

FIGURE 5-6. MODEL DERIVED EXPECTED UTILITY INDIFFERENCE POINT VS. SUBJECTS STATED INDIFFERENCE POINT FOR ICEBERG AND SOMETHING SENSORS

A measure of the variability of the utilities is the Utility Matrix Difference (UMD) score. This measure is computed as follows:

$$\mathrm{UMD}(t_1, t_2) = \sum_{k,i} \left| {}_kU^+_{it_2} - {}_kU^+_{it_1} \right| + \sum_{k,i} \left| {}_kU^-_{it_2} - {}_kU^-_{it_1} \right| \qquad (5\text{-}1)$$

The UMD score is a measure of the variability of the utility values from cycle t_1 to t_2.[1] In the following analysis a global measure is used, which summarizes the variability of the utilities for the entire session. The session UMD score is the sum of the single-cycle UMD scores from the start of the session, t_o, to the end of the session, t_e. It is defined as:

$$\mathrm{SUMD} = \sum_{t=t_o}^{t_e-1} \mathrm{UMD}(t, t+1) \qquad (5\text{-}2)$$

Results. The SUMD scores for the fourth session, the session in which each group received different aiding treatment, are shown in Table 5-3 for each subject. Group 3, the group which received aiding and an explanation of how the aiding was derived from their own behavior, had the lowest mean SUMD score and range of values. This indicates that the operator's decisions were extremely consistent with the model derived decisions. Group 2, the group which received aiding but did not receive any explanation of how the aiding was derived, had the same mean degree of inconsistency as Group 1, the group which did not receive any aiding at all. Group 2, however, had a much wider range of scores. A specific comparison between Group 3 and the other two groups (utilizing the non-parametric Mann-Whitney U test) indicated the observed difference was statistically reliable beyond the 0.05 level.

---

[1] In the present task, time advances in discrete steps, one step to a decision cycle.

TABLE 5-3

UTILITY VARIATION DURING FOURTH SESSION (N=9)

SESSION UTILITY MATRIX DIFFERENCE SCORE (SUMD)

	CONTROL GROUP (G1)	AIDING ALONE (G2)	AIDING AND INDOCTRINATION (G3)
SUBJECT 1	81	13	0
SUBJECT 2	90	119	14
SUBJECT 3	138	199	25
MEAN	103	110	13
RANGE	57	186	25

G3 is significantly different from G1 and G2 combined.
(P=0.048 Mann-Whitney U-Test; n_1=3, n_2-6)

Table 5-4 shows the number of status reports filed in the fourth session. This score is a measure of the performance rate during the fourth session since the session length was the same for each subject. Group 2 and 3 appear to have slightly higher performance scores than Group 1. This difference was marginally significant at the 0.08 level.

Table 5-5 shows the UMD score/status report ratio for each subject. This score represents a normalized utility variation for the fourth session. As expected, these data are similar to that of Table 5-3. The Mann-Whitney U-test indicated that Group 3 is reliably different from Group 1 and Group 2 beyond the 0.025 level.

Discussion. The major finding of the investigation was that decision aiding with apparent control leads to more consistent decision behavior. Similar results were reported by Hanes and Gebhard (1966) who used Navy commanders in a realistic task simulation. There was also some evidence that Group 2, which received aiding but had little understanding of the aiding process or control over it, had more extreme responses to the aiding. This is consistent with the postulation by Halpin, et al (1973), that a lack of knowledge and understanding will lead to a more extreme -- much higher or much lower -- evaluation and weighting of aiding suggestions than is appropriate.

The better performance (in terms of the number of status reports filed) shown by the groups given aiding appeared to be due both to faster sensor placement and to less vacillation when events in the intelligence report fell close to the subject's indifference points. This improvement in performance was a byproduct of the experiment rather than a direct objective. It is clear that if increasing throughput were a primary goal, considerable improvement could be obtained. One way would be to call the operator's attention to the need for fast decisions. This undoubtedly would accentuate the effect of aiding on the speed of decision making.

TABLE 5-4

PERFORMANCE RATE DURING FOURTH SESSION (N=9)

NUMBER OF STATUS REPORTS FILED

	CONTROL GROUP (G1)	AIDING ALONE (G2)	AIDING AND INDOCTRINATION (G3)
SUBJECT 1	29	43	34
SUBJECT 2	35	42	51
SUBJECT 3	32	31	39
MEAN	32.0	38.7	41.3
RANGE	6	12	17

G1 is marginally different from G2 and G3 combined.
($p=0.08$; Mann-Whitney U-Test; $N_1=3$, $n_2=6$)

TABLE 5-5

NORMALIZED UTILITY VARIATION DURING FOURTH SESSION (N=9)

SUMD SCORE/STATUS REPORTS FILED

	CONTROL GROUP (G1)	AIDING ALONE (G2)	AIDING AND INDOCTRINATION (G3)
SUBJECT 1	2.79	0.30	0.00
SUBJECT 2	2.57	2.83	0.28
SUBJECT 3	4.31	6.42	0.64
MEAN	3.22	3.18	0.31
RANGE	1.74	6.12	0.64

G3 is significantly different from G1 and G2 combined.
(P=.024; Mann-Whitney U-Test; $n_1=3$, $n_2=6$)

Another way would be to partially or completely automate the decision process. The implementation of a "reject" system as suggested by Hanes and Gebhard (1966) would accomplish this. The recommended decisions would be accepted automatically, unless explicitly rejected by the operator. Such decision making by default would clearly improve the speed of decision making. Comparison of the indoctrinated and unindoctrinated aiding groups suggests that the operator's knowing that the automatic decisions are essentially his own would make this form of rapid decision making both acceptable and effective.

The acceptance of decision aiding by the operator has the effect of making his decisions more consistent with a normative EU model of his behavior. This apparently provides a mechanism which diminishes the effect of some biases that are incompatible with other beliefs held by the operator. It does not eliminate biases since they may be incorporated into the utility estimates. Also, acceptance of aiding helps the operator to overcome some of his limitations of memory and ability to categorize or group events and outcomes.

Becoming more consistent with the normative model allows the operator to more readily evaluate the adequacy of his strategies and his utility structure. For example, it facilitates the systematic evaluation of those situations which are and are not adequately handled by the model and by the operator's utility structure. It allows the operator to discover more readily the limitations of his approach to the task and the possible need for changing his value structure in order to improve his performance. In the absence of consistent behavior, it is extremely difficult for him to evaluate the adequacy of his approach because of limitations on his memory and his ability to group outcomes.

6. REFERENCES

1. Freedy, A., D. May, R. Weisbrod, and G. Weltman, "Adaptive Computer Aiding in Dynamic Decision Processes: Part II, Scenario Generation by Elicited Expert Probabilities," Technical Report #PTR-1016-74-5(II), Perceptronics, Inc., Woodland Hills, California, ONR, N00014-73-C-0286/NR 196-128, Order No. 2347, ARPA, Dept. of Defense, May 1, 1974.

2. Freedy, A., Weisbrod, R., Davis, K., May, D., and Weltman, G. "Adaptive Computer Aiding in Dynamic Decision Processes: Part I, Adaptive Decision Models and Dynamic Utility Estimation," Technical Report #PTR-1016-74-5(I), Perceptronics, Inc., Woodland Hills, California, ONR, N00014-73-C-0286/NR 196-128, Order No. 2347, ARPA, Dept. of Defense, May 1, 1974.

3. Hanes, R.M. and J.W. Gebhard, "The Computer's Role in Command Decision," U.S. Naval Institute Proceedings, September, 1966.

4. Halpin, S.M., J.A. Thornberry and S. Streufert, "The Credibility of Computer Estimates in a Simple Decision Making Task," Technical Report No. 5, Purdue University, ONR, N00014-67-0226-0030/NR 177-946, Org. Effectiveness Res. Program (452), Office of Naval Research, January, 1973.

5. Kneppreth, Norwood P., David H. Gustafson, Edgar M. Johnson, and Richard P. Leifer, "Techniques for the Assessment of Worth," Technical Paper 254, DAHC 19-72-C-0026, U.S. Army Research Institute for the Behavioral and Social Sciences, December, 1973.

6. Peterson, Cameron R., "Judgements of Probability and Utility for Decision-Making," Report 037230-1-A, Engineering Psychology Laboratory, University of Michigan, ONR, N00014-67-A-0181-0034, Engineering Psychology Programs, Office of Naval Research, September, 1971.