AD-A014 218

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES:
AN EXPERIMENTAL STUDY OF AIDING EFFECTIVENESS

Kent B. Davis, et al

Perceptronics, Incorporated

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

1 May 1975

ADA014218

ADAPTIVE COMPUTER AIDING

IN DYNAMIC DECISION PROCESSES:

AN EXPERIMENTAL STUDY OF AIDING EFFECTIVENESS

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

DDC
AUG 29 1975
RECEIVED
B

PERCEPTRONICS

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

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES:
AN EXPER:MENTAL STUDY OF AIDING EFFECTIVENESS

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

PERCEPTRONICS

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

i.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER PTR 1016-75-5	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES: AN EXPERIMENTAL STUDY OF AIDING EFFECTIVENESS		5. TYPE OF REPORT & PERIOD COVERED Semi-Annual Technical Report 10-1-74 to 4-31-75
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Kent B. Davis, Ricnard L. Weisbrod, 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 May 1, 1975
		13. NUMBER OF PAGES 56
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	Man-Machine Interaction
Control	Man-Computer Interaction
Computer Aided Decision	Utility Estimation
Adaptive Learning Systems	

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

This report describes an experimental study which was performed to evaluate 1) the effectiveness of adaptive decision aiding for improving decision maker (DM) performance and 2) the consistency of decision behavior in a situation where the DM can develop a viable strategy. Twelve Naval Reserve non-commissioned officers gathered intelligence about a simulated fishing fleet moving about an expanse of ocean and reported its status.

DD FORM 1 JAN 73 **1473**   EDITION OF 1 NOV 65 IS OBSOLETE

1a

20.

Nine different types of sensors with overlapping capabilities and varying reliabilities and costs gave the operator a realistic set of decision alternatives. As he performed the task, the ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System tracked the operator's decision behavior and, by means of adaptive pattern classification techniques, estimated his utilities. These utility estimates are a measure of the relative worth to the operator of the information outcomes of each sensor placement.

The subjects were divided into two equal groups. After three training sessions, the experimental group was given aiding in the form of sensor deployment recommendations. The recommendations were based on a maximum expected utility criterion which employed estimates of the subject's own utilities. The control group received no aiding.

Using deviation from maximum expected utility as a performance measure, it was found that the aided experimental subjects performed significantly closer to the maximum EU level than did the unaided control subjects. It was also found that the unaided group showed significantly more within-group variability in their performance than the aided group.

The report also provides further elaboration of the significance of the utilities used in the ADDAM model. In addition, it describes a number of modifications and improvements which have been made to the ADDAM system as a result of analysis of earlier experiments. These changes include several new sensor types and sensor deployment commands introduced to give the operator a more realistic set of decision alternatives. Other changes include a more responsive utility training algorithm and the introduction of performance measures.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

1. SUMMARY

1.1 Purpose

A large number of present-day military systems require operations personnel to make rapid and complex sequential decisions. These include systems for:

    (1) ASW Command and Control

    (2) Shipboard Tactical Operations

    (3) Aircraft Combat and ECM

    (4) Ground/Air Resource Allocation

    (5) Remotely Piloted Vehicles

As systems continue to increase in size, sophistication, and speed, the costs of sub-optimal decisions grow correspondingly greater.

The research described here is concerned with computer methods for aiding the operator to make better and more timely decisions. It uses as a representative decision situation a simulated surveillance task, in which the operator uses the sensor resources at his command to monitor and to report the movements of a mobile fishing fleet. The decisions made by the operator are quite similar to those required in ASW Operations, in Tactical Resource Allocation, and in Remote Aircraft Control, to mention a number of actual tasks which have already been mapped into our research structure. The computer system developed for decision aiding in the present program is termed ADDAM. Its design has been described in a previous project report (Freedy, Weisbrod, Davis, May, and Weltman, 1974).

In brief, the ADDAM (Adaptive Dynamic Decision Aiding Mechanism) system uses an adaptive utility estimation program to determine an operator's value

structure in a sequential decision task, and supplies decision aiding based on the operator's own preferences. The estimated utilities provide direct, on-line measures of the operator's decision making behavior. A pilot study (Weisbrod, Davis, Freedy, and Weltman, 1974) indicated that the system measures stabilized rapidly, and that aiding in the form of recommended decisions significantly improved operator decision making performance. The present study was designed to verify and extend these findings by examining a more complex task and a larger subject group.

1.2 Adaptive Decision Aiding

The primary function of ADDAM is not simply to model the decision maker's behavior, but to provide a basis for decision aiding. Once ADDAM has learned the operator's values, it applies them in several ways, e.g., to suggest decisions which maximize his return, to point out inconsistencies in strategy, etc. Because the decision model is adaptive, decision aiding establishes a complex symbiotic relationship between the operator and ADDAM. ADDAM adapts to the human pattern of behavior and, in turn, provides decision aiding which may cause the human to modify his behavior.

1.3 Experimental Study

Twelve male subjects were recruited from nearby Naval Reserve units. Each subject was given three individual training sessions with the fishing fleet task, during which he practiced information acquisition and analysis through the placement of remote sensors. In a fourth session, subjects were randomly divided into control and experimental groups of 6 subjects each. The experimental group received decision aiding while the control group did not. The decision aiding consisted of recommended sensor allocations, and was based on utilities estimated previously for each individual.

1.4   Experimental Results

Consistency. The operator's mean deviation from maximum expected utility (DEVMAXEU) during the course of the test session was the primary measure of decision performance. As a group, the aided operators were markedly more consistent in their decision making performance, while the control group exhibited extremes of behavior. An F-test conducted on the differences in group variance was significant at the .001 level.

Effectiveness. The aided operators performed much closer to maximum expected utility. In fact, some members of the control group produced gross deviations from optimum performance. Because of the differences in group variance, a log transformation was applied to the DEVMAXEU scores. A subsequent F-test on a log transformation of the group means was significant at the .05 level.

Decision Output. The number of task cycles completed during the test session is a measure of decision efficiency. The aided group had a significantly greater mean output (F-test, P < .025) than the unaided control group.

Utility Convergence. Data was obtained to determine the rate at which the utility estimation program responded to the operator's decision making preferences. There appeared to be two stages in machine adaptation: 1) a rapid stage, in which the major portion of adaptation is made; and 2) a gradual stage, in which minor adjustments are accomplished. In a typical case, major adaptation was completed in only five decision cycles.

1.5   Conclusions

The present findings replicated the results of the pilot study.
To date, twenty-one operators have been examined in an experimental
context and the outcome has been consistent. Accordingly, the following
conclusions can be stated with some degree of certainty:

(1) The ADDAM program adaptively estimates operator utilities in
     realistic decision making situations.

(2) Utility estimation is consistent over subsets of the total
     outcome set.

(3) Utility estimation rapidly stabilizes for consistent operator
     decision behavior.

(4) Decision recommendations based on adaptive utility estimates
     are well accepted by experienced operators.

(5) Availability of individualized recommendations markedly
     improves decision making performance by a) allowing the
     individual operator to maintain near maximum expected
     utility; and b) reducing variability among different
     operators.

Of particular importance for the use of adaptive aiding is
conclusion (2), which indicates that a large set of utilities can be
trained in trials involving only a small number at a time, and conclusion
(3), which indicates that utilities may be estimated in a reasonably
short time period. Our on-going analysis of military situations indicates
that for many of them, the decision making requirements closely match

those of the ADDAM system. Thus it appears that practical adaptive
computer aiding systems are feasible and may offer important improvements
in decision making effectiveness. According, our subsequent work will
include a closer examination of application areas.

2. INTRODUCTION

2.1   Overview

The ADDAM System is based on a principle of on-line acquisition of operator decision strategies by means of direct observation of his behavior. This principle requires on-line observation of operator decision made in response to real world probability data in order to computationally infer his value structure and provide parameters for a behavioral model. This approach, termed "on-line model matching", has been practiced in adaptive control for some time (Margolis and Leondes, 1960). It is intended for complex situations where direct analytical techniques are cumbersome and inefficient. One major advantage of this technique is that model parameters are continuously checked against the real system and adjusted to track changes in system behavior.

In applying the on-line model matching approach to dynamic decision situations, three major elements are required: (1) an adequate theoretical structure for modeling the decision process, (2) model learning and parameter estimation techniques, and (3) a computational capability for real-time observation and modeling of decision processes. ADDAM combines these three functional requirements in a viable system.

The ADDAM system uses an adaptive expected utility (EU) model as a paradigm for operator behavior. Since the EU equations have the same mathematical form as a linear discriminant function, pattern recognition techniques adaptively estimate model parameters. The probabilities of the decision outcomes correspond to the input patterns, the decision maker's (DM's) utilities for these outcomes correspond to the weight parameters, and the DM's decisions correspond to the correct classification of the input pattern. In a pattern classifier, the weight parameters are adjusted until

the pattern recognizer can correctly classify the input patterns. In the
ADDAM system, an on-line computer system operating in parallel with the
human decision maker predicts what the DM will do according to the EU model.
The computer then compares the prediction with the DM's actual decision and
trains the utilities in a "behavioral" manner. Incorrect predictions are
punished and correct predictions are rewarded by an error correcting
algorithm which adjusts the utilities in a large decision matrix. Thus
ADDAM learns the DM's values for decision outcomes and uses them to make
increasingly accurate predictions of his behavior. And when the DM changes
his behavior over time, ADDAM responds adaptively to these changes. Detailed
technical descriptions of the model and the ADDAM system structure are pro-
vided in earlier technical reports (Freedy, Weisbrod, Davis, May, and
Weltman, 1974; Weisbrod, Davis, Freedy, and Weltman, 1974). A brief review
is found in Chapter 3.

The primary function of ADDAM is not simply to model the decision
maker's behavior but to provide a basis for decision aiding. Once ADDAM has
learned the operator's values, it is possible to use them to aid him in
several ways: (a) suggesting decisions which optimize his values, (b) warning
him when he makes suboptimal decisions, (c) calling his attention to critical
events, and (d) providing a basis for comparing his value structure to
organizational standards.

Because the decision model is adaptive, model based decision aiding
establishes a complex symbiotic relationship between the operator and
ADDAM. ADDAM adapts to the human operator's pattern of behavior and, in
turn, provides decision aiding which may cause the human to modify his
behavior.

2.2   Decision Aiding in an Operational Environment

This section briefly reviews a military operational environment under which decision aiding may be required and explores the applicability of decision aiding based on the ADDAM concept. Military operational missions can be divided into three phases, each with their own decision requirements. The first phase involves planning. The courses of action which may be pursued during the mission are selected. These alternatives are based on mission objectives and externally imposed constraints. Examples of constraints include available hardware, established procedures, manpower limitations, and political factors.

A plan is a set of selected actions that define the mission. The decision processes which lead to the selection of these actions involve detailed analysis and evaluation of decision alternatives and their possible outcomes. Because the outcomes of each alternative can be defined only probabilistically, the decisions are made under risk.

The second phase of an operational mission is the execution phase. This phase involves continuous decision processes in a dynamic environment. In tactical operations, factors such as enemy movements and counter movements, weather changes, intelligence data, etc. must be evaluated and the plan modified as the environment and the decision maker's awareness of it change. The results of previous decisions usually establish constraints which affect subsequent decision alternatives.

The final phase of an operational mission is the evaluation phase. The decisions made in the planning and execution phases are reviewed and evaluated. Recommendations are made to modify the planning procedures and decision strategies for future situations.

Decisions made in such an operational environment can be classified according to type. Examples include resource allocation, logistic decisions, configuration of resources, configuration of forces, and intelligence decisions. Each type of decision can also vary from one military context to another (e.g., anti-submarine warfare).

The decision aiding techniques which have potential application in each type of decision situation depend on the requirements of that particular situation. Classes of techniques include the following: (1) Data Organization and Display. This type of aiding may involve data files which are organized for presentation to the decision maker in the form of reports, etc. (2) Decision Procedures. Decision procedures attempt to establish a prescribed course of action which will insure high quality decision making. These procedures may involve guidelines for selecting relevant alternatives, analyzing their outcomes and their probabilities, and assessing new information. (3) Mathematical Optimization. This form of decision aiding is applicable to situations where it is possible to establish objective decision criteria based on a mathematical measure for assessing the expected gain from each alternative course of action. Mathematical optimization techniques such as linear programming, dynamic programming, and other techniques derived from control theory, have been successfully applied to engineering design and system optimization decisions. (4) Decision Analysis. Decision analysis establishes a prescriptive procedure for decision action selection and optimal decision making on the basis of normative decision rules for rational choice (Howard, 1968; Brown, Hoblitzell, Peterson and Ulvila, 1974). Decision analysis also establishes a procedure for decomposing decision and subsequently maximizing the expected gain of the decision maker.

2.3   Aiding by Means of Decision Analysis

Applying decision analysis to decision aiding involves seven basic
steps. These steps cover the relevant aspects of defining decision choices
and parameters required to establish a measureable criterion of optimal
choice (Payne, Miller, Ronney, 1974):

1. Identification of pertinent information.
2. Definition of alternatives.
3. Definition of structure for related data parameters, events
    and alternatives.
4. Characterization of uncertainty of continuous parameters.
5. Estimation of event probabilities.
6. Transformation of multi-attribute measures into a single
    utility for each possible outcome.
7. Selection of the best alternative through normative evaluation
    criteria.

The basic components of the criteria are probabilities and utilities.
It is necessary to determine the probabilities of alternative decision
outcomes and assess the utilities that the decision maker has for these
outcomes. Probabilities can normally be estimated by objective measurement
or from prior probabilities elicited from experts. Then the prior proba-
bilities can be aggregated, using Bayesian or probabilistic information
processing (Edwards, 1962; Kelly and Peterson, 1971), to obtain posterior
probability estimates. These techniques also provide a mechanism for up-
dating the probabilities as new data becomes available.

Reliable methods for quantitative assessment of utilities are a
major area of difficulty. A number of techniques have been suggested and
used, both in the research literature (Kneppreth, Gustafson, Leifer, and

2-5

Johnson, 1974), and in the emerging discipline of applied decision analysis (Brown, Hoblitzell, Peterson and Ulvila, 1974). The problem of assessing utilities has become especially acute in recent years because of the growing interest in quantitative decision analysis.

2.4   Methods of Quantitative Utility Assessment

A large number of techniques for utility assessment have been suggested. These may be classified according to the measurement and computational processes used to estimate them. Comprehensive reviews of utility assessment techniques have been prepared (e.g., Kneppreth, Gustafson, Leifer, and Johnson, 1974; Fishburn, 1967). Three major classes of utility assessment techniques are briefly reviewed here:

(1) Elicitation of utilities through direct judgement. An analyst asks the DM directly to give his value for each decision outcome. These values are normally measured as point values for a particular outcome. These values can be obtained directly using a wholistic approach (Beach, 1973). However, since outcomes usually have several attributes they are often decomposed into single attribute outcomes. The single attribute utilities thus elicited are then combined linearly, to yield the DM's utilities for the more complex outcomes.

(2) Inference from behavior in simple gambles and other decision games. This technique requires the DM to respond to a series of simple gambles or decision games which usually involve financial reward. The DM's choices form a data base from which his utilities are inferred, usually by indifference techniques. These techniques have been used by a number of investigators (e.g., Tversky, 1967), but are mainly constrained to laboratory/ research settings.

(3) Dynamic estimation through decision observation. This approach
calls for direct observation of decision behavior in real-world
or simulated real-world contexts. The primary example of this
approach is the ADDAM System (Freedy, Weisbrod, Davis, May, and
Weltman, 1974). Other attempts to use this approach are more
concerned with modeling the decision maker's gross behavior than
with determining his utilities. One good example is the boot-
strapping techniques of Dawes (1970) which use a brute force
linear model.

A comparison of the positive attributes of utility assessment tech-
niques is illustrated in Table 2-1. The advantages of the Dynamic Observa-
tion technique are as follows: (1) Utilities are estimated non-verbally,
without the need for a skilled analyst highly trained in utility estimation
techniques. Indeed, the DM need not be aware that his utilities are being
assessed. Utilities can be estimated rapidly and the technique is not
limited by the number of possible decision outcomes. (2) The utilities are
measured on a common scale and are combinatory. (3) The utility assessment
technique responds adaptively to changes in values and the utilities are
automatically validated by direct comparison with the DM's real-world
behavior. These advantages have important implications for decision aiding.

Dynamic utility estimation provides a basis for real-time decision
aiding. The utility estimates are based on the DM's current behavior rather
than on static utilities estimated in an off-line context. Thus, in situ-
ations where the DM's values change over time, the aiding he receives is
based on his current utilities.

2.5   Experimental Validation of Aiding

The primary objective of the current experimental study is to provide direct evidence of decision aiding and to evaluate the effectiveness of the adaptive aiding system. A previous study (Weisbrod, Davis, Freedy, and Weltman, 1974) using three college students per treatment group, found ADDAM 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 (95%) in predicting the operator's decisions. The adaptive decision model was also found to be sensitive to individual differences in decision strategies.

Decision aiding, presented to subjects who had knowledge of its adaptive nature, resulted in a higher degree of consistency with the normative decision model. Aiding without such knowledge appeared to accentuate individual differences in behavior. Decision aiding also appeared to improve the decision throughput of the subjects by allowing them to place sensors more quickly and by reducing the amount of vacillation near the indifference points.

The current experimental study differs from the previous one in several important aspects: (1) twelve subjects, recruited from Naval Reserve units were used; (2) the decision space is much more complex; (3) a more sensitive training algorithm is used to decrease the amount of time necessary to train the model; (4) more of the subject's behavior was sampled than in the previous study; and (5) deviation from optimum expected utility is used as a performance standard for analysis.

The current study uses a one-way experimental design with two treatment levels. The twelve subjects were divided between a no-aiding

TABLE 2-1. COMPARISON OF UTILITY ASSESSMENT TECHNIQUES

POSITIVE ATTRIBUTE	DIRECT ELICITATION	GAMES AND GAMBLE BEHAVIOR	DYNAMIC OBSERVATION
NON-VERBAL	NO	NO	YES
NUMBER OF OUTCOMES	UNLIMITED	LOW (TWO)	UNLIMITED
PRIOR TRAINING OF DECISION MAKER	MODERATE	EXTENSIVE	NONE
SPEED	FAST	SLOW	FAST
SKILLED ANALYST REQUIRED	YES	YES	NO
REAL WORLD VALIDATION	NO	NO	YES
COMMON UNITS	AFTER WEIGHTING	YES	YES
COMBINATORY	AFTER ANALYSIS	YES	YES
ADAPTIVE TO VALUE CHANGES	NO	NO	YES

(control) group and an aiding (experimental) group. 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. During the fourth session, the experimental group received decision aiding in the form of sensor deployment recommendations derived from the adaptive decision model, while the control group continued to perform the task as before. Chapter 4 presents a more detailed discussion of the experiment.

3. ADDAM SYSTEM

The ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System provides a vehicle for research on dynamic decision making, adaptive decision models, dynamic utility estimation, and man/computer decision making. ADDAM consists of a system for simulating a dynamic decision task, an adaptive decision model based on dynamic utility estimation, and mechanisms for man/computer interaction and decision aiding.

3.1   Fishing Fleet Simulation

The fishing fleet simulation provides a continuous on-line decision task. Basically, the operator must track the elements of a simulated fishing fleet as it moves about in an expanse of ocean and report their status. To perform the task the operator deploys sensors of varying reliabilities and costs.

Fleet Elements. The elements of the fleet consist of a trawler which periodically deploys its net, and an iceberg. Both move about on an ocean which is represented by a 5 by 5 grid (see Figure 3-1). The elements move just prior to the start of each task cycle (described below). They may remain stationary or move to an adjacent grid location to the North, South, East, or West. The elements cannot move off the grid.

Figure 3-2 illustrates a typical sequence of movements. In this example, the iceberg starts (Figure 3-2a) at the upper edge of the board and begins to move south (down) one square at a time. At the start of the fifth cycle (Figure 3-2e) it begins to move east. The trawler begins near the bottom of the board, moves one square to the east, deploys its net, retracts the net, and finally heads north.
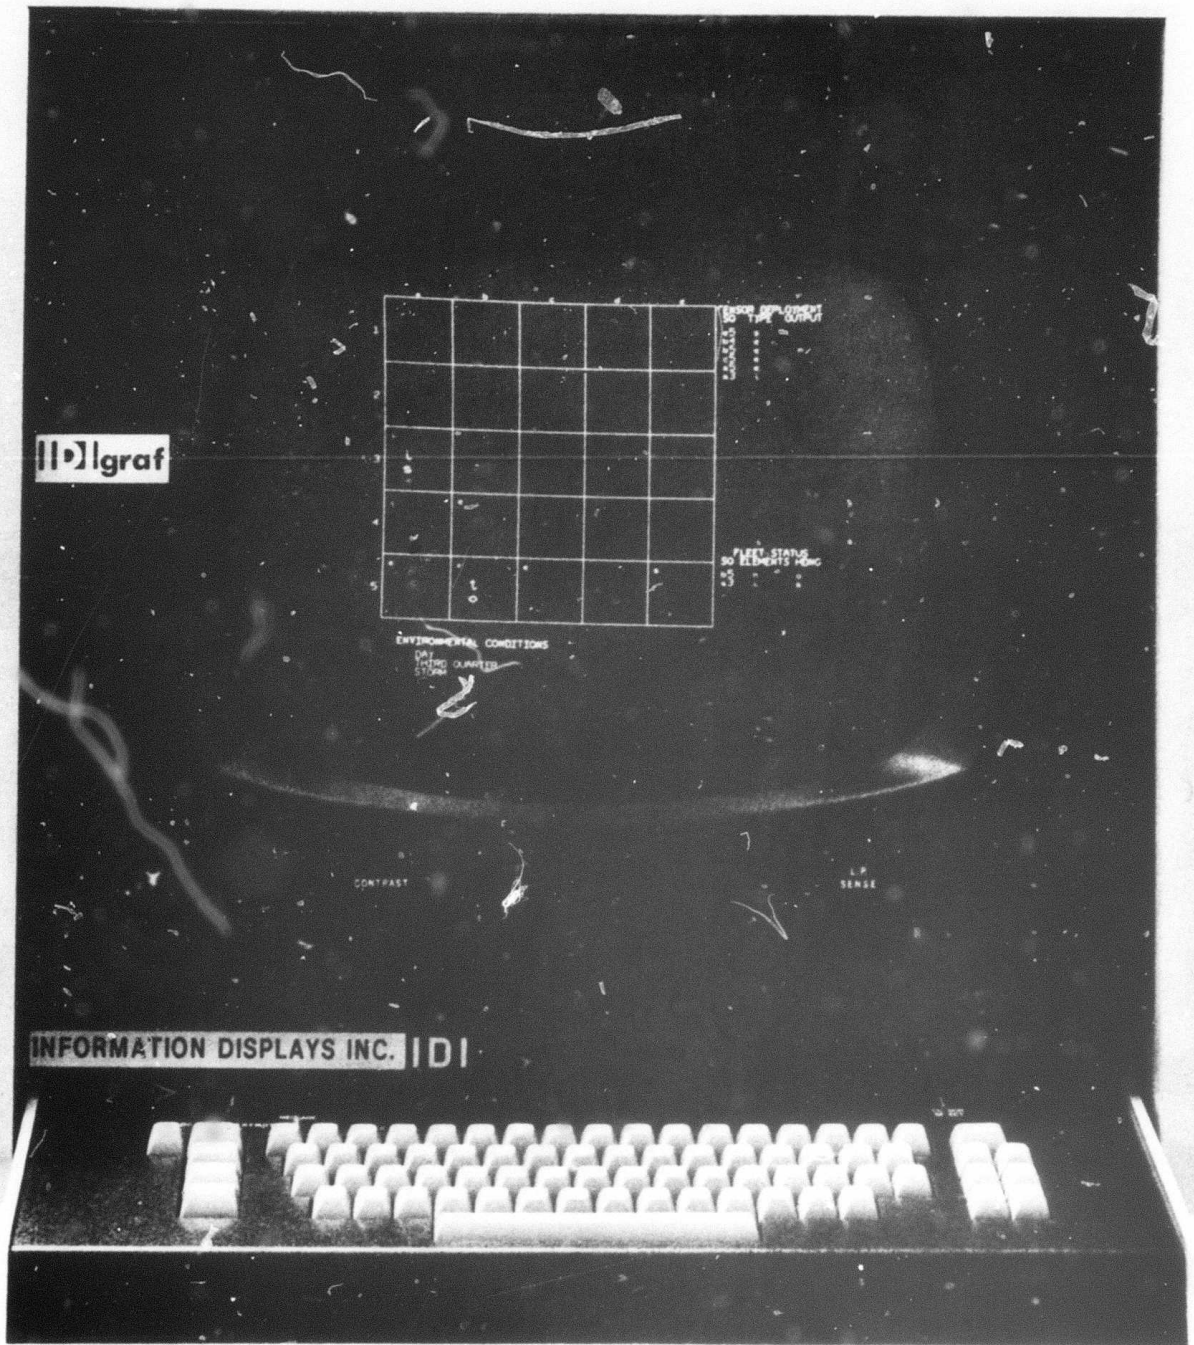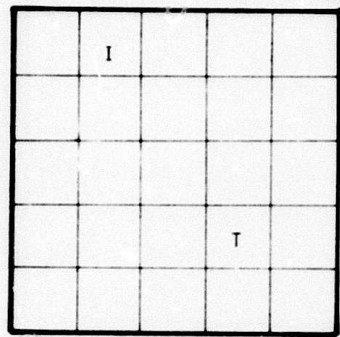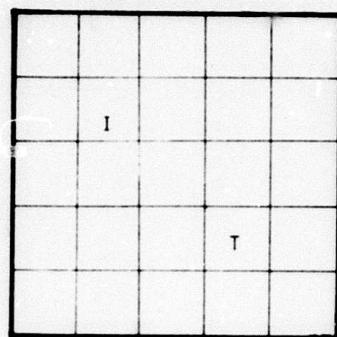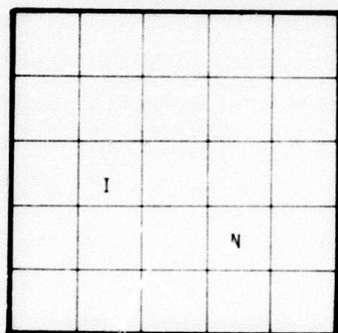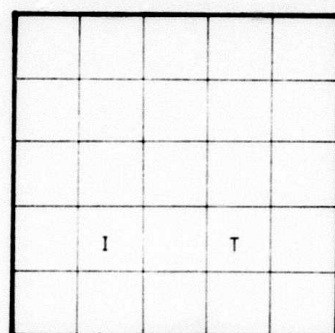
FIGURE 3-1.   OPERATOR INFORMATION DISPLAY
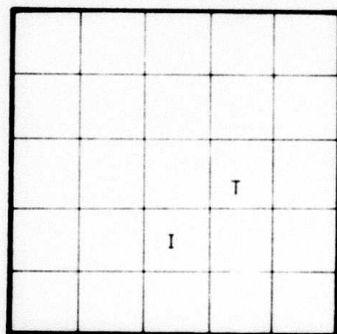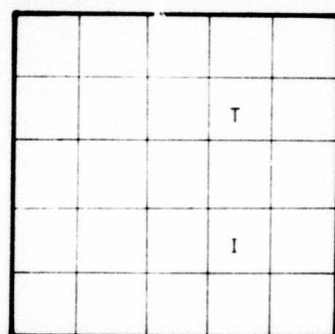
FIGURE 3-2. TYPICAL FISHING FLEET SCENARIO

Sensors. The operator cannot observe the movements of the fleet elements directly. His only access to the environment is through sensors which he deploys at selected grid locations. These sensors differ in their abilities to detect different types of objects and in their reliabilities. An "iceberg" sensor, for example, can only detect icebergs, while an "everything" sensor can detect any type of object with high reliability. An everything sensor, however, is very costly to use. Table 3-1 summarizes the properties of the eight different types of sensors available.

Sensors can detect objects only at the locations where they have been deployed. They cannot detect objects at adjacent locations. The operator has an unlimited number of each type of sensor, but he cannot deploy more than one sensor at each grid location. All sensors are removed automatically at the end of each decision cycle and must be redeployed at the start of the next cycle.

In the current fishing fleet simulation, the only kind of error a sensor can make is to report a false alarm. Thus, a sensor might falsely report the presence of a fleet element, but it will never fail to report an object which is actually present. The rate at which false alarms occur is a property of the sensor type and is listed in Table 3-1.

Operator Task. The operator's task is to monitor the movements of the fishing fleet elements and to report their locations. To perform this task the operator deploys sensors, reads their outputs, reports the status of the elements, and receives an intelligence report which he uses to make his next round of decisions. In some experimental contexts, he also receives decision aiding generated by the ADDAM system. This decision task sequence is illustrated in Figure 3-3. The following paragraphs give a more detailed account of what the operator sees and does.

TABLE 3-1. SENSOR CHARACTERISTICS

TYPE	OBJECT SENSITIVITY	FALSE ALARM RATE	COST
T	Trawler	5%	$2.50
V1	Trawler & Net	20%	$2.50
V2	Trawler & Net	35%	$1.50
N	Net	5%	$2.50
I1	Iceberg	10%	$2.50
I2	Iceberg	30%	$1.25
S	Something	40%	$1.00
E	Everything	5%	$7.00

The operator sits at the IDIgraf graphics display terminal and is confronted by a representation of the 5 by 5 environment grid. The input cursor is positioned in the Sensor Deployment Field at the upper right hand corner of the display. Since the operator does not have any indication of where the trawler and iceberg are, he enters Search mode by typing in the appropriate command on the keyboard. Search mode allows the operator to deploy everything sensors over the entire board. Once the operator has found the objects, sensor deployments will be made selectively according to some strategy.

When the operator has finished deploying sensors he types "f" and the sensors appear in the upper left hand corner of the squares in which they were deployed. There is a momentary pause, while experimental data is printed out, and then the sensors are activated. Those which have a positive response begin to blink and the sensor output appears in the output column to the right of the Sensor Deployment Field entry, as illustrated in Figure 3-1. The figure shows that the everything sensor at location b5 has detected a trawler with its net deployed. This is indicated by the "n" in the output column.

The operator reads the sensor outputs and decides where the objects are. The display input cursor is automatically positioned to the fleet status field at the lower right of the display and the operator enters his status decisions. When he finishes entering his status report he types "f". The report is entered onto the board as illustrated in squares a3 and b5 of Figure 3-1, and is transmitted to the intelligence report generator.

The intelligence report is transmitted to the operator via the teletype attached to the system. The intelligence report represents an expert analysis of what the fleet elements would do if the operator's

3-6
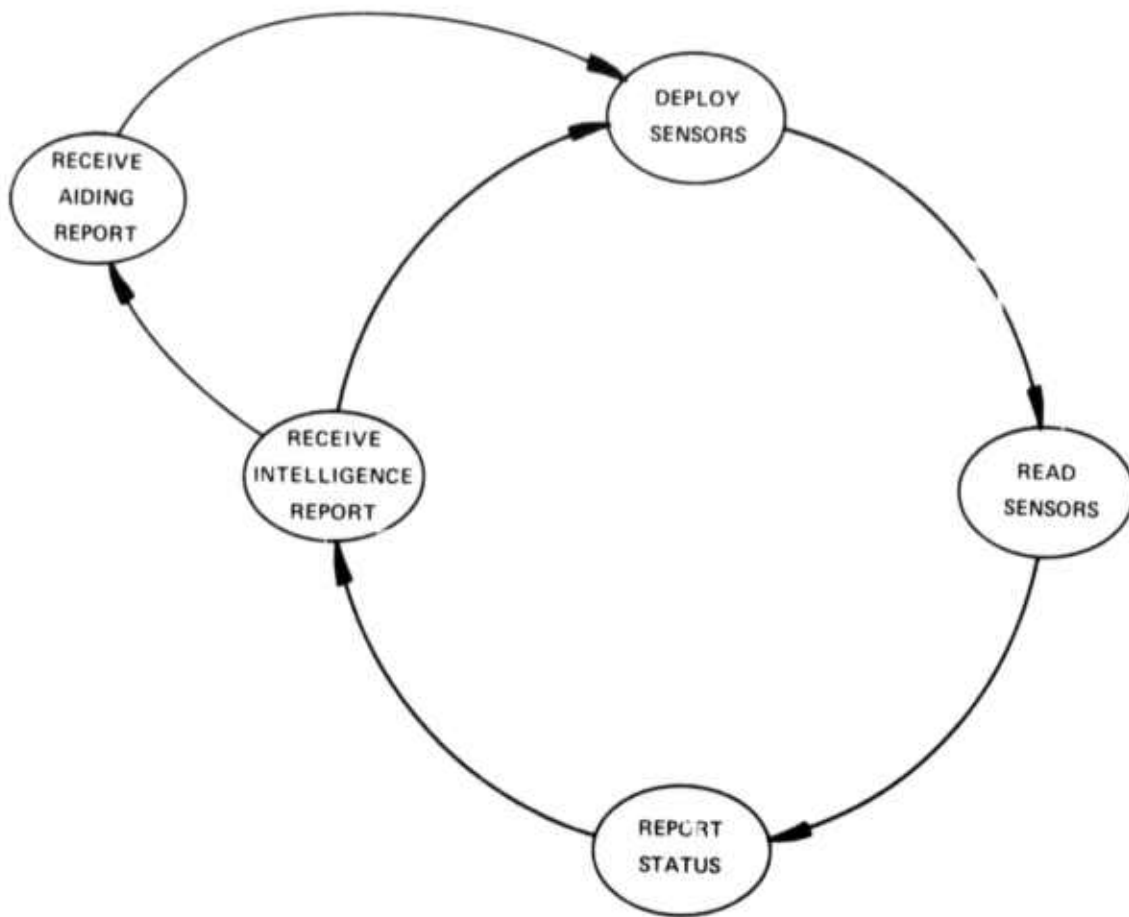
FIGURE 3-3.   DECISION TASK SEQUENCE

status report were correct. The report lists the probabilities that each
type of object will be at each board location. Locations which have zero
probability of an object are not listed. A typical intelligence report is
illustrated in Figure 3-4.

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 3-4. INTELLIGENCE ANALYSIS REPORT

The transmission of the intelligence report to the operator is the
final step in the decision task sequence. In experiments involving
decision aiding there is an additional step. After receiving the intelli-
gence report, the operator also receives decision aiding in the form of
sensor deployment recommendations. These recommendations appear as a
checklist in the sensor deployment field of the graphics display terminal.
Prior to entering his own sensor decisions, the operator accepts, rejects,
or modifies the suggestions, one by one. He then adds his own decisions
to the end of the list.

3.2   Functional Description

The functional organization of the ADDAM system is illustrated in Figure 3-5. The environment generator probabilistically generates the dynamic decision environment on the basis of expert probabilities and an organization structure specified by the experimenter. The environment, as seen through sensors deployed by the operator, is displayed on the 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.

Scenario Generator. The scenario generator simulates the movements of the fishing fleet elements and the responses of the operator's sensors. It does this by means of a unique application of Bayesian probabilistic information processing (PIP) techniques. Instead of aggregating expert opinions to estimate the probabilities of complex events in a real world, these probabilities are used to simulate the real world.

The technique for simulating the fishing fleet is based on a transition matrix which relates the current state vector to a set of state transformation operators. The components of the transition matrix are the conditional probabilities of each state transformation operator (or rule

which dynamically changes the environment), given the value of each state vector variable. These conditional probabilities are estimated by "experts" on the behavior of the environment being simulated.

The next step is to compute the conditional probabilities of each state transformation operator, given the current state vector. The actual state transformation operators applied to the current state vector are chosen on the basis of these probabilities by means of a Monte Carlo selection procedure. The state transformation operators are then executed to obtain a new state vector. A more detailed discussion of the scenario generation technique is found in Freedy, May, Weisbrod, and Weltman, 1974.

Intelligence Report Generator. The intelligence report generator is essentially a simulated Bayesian probability estimator. An expert intelligence analyst is simulated by using the status of the fishing fleet (reported by the operator) as the state of the real world. The environment generator's expert conditional probability matrix is aggregated in a conventional PIP manner to obtain the prior probabilities of the next state of the environment. These probabilities are the ones which would actually be used to generate the next state if the reported status accurately reflected the current state. A sample intelligence report is illustrated in Figure 3-4.

Utility Estimation. The dynamic utility estimation technique is based on the principle of a trainable multi-category pattern classifier. The utility estimator observes the operator's choices among the available decision options and attempts to classify the patterns of event probabilities (i.e., the intelligence report) by means of an expected utility discriminant function. Decisions predicted by this model are compared with the operator's actual decisions. Whenever the model's classifications are incorrect, the
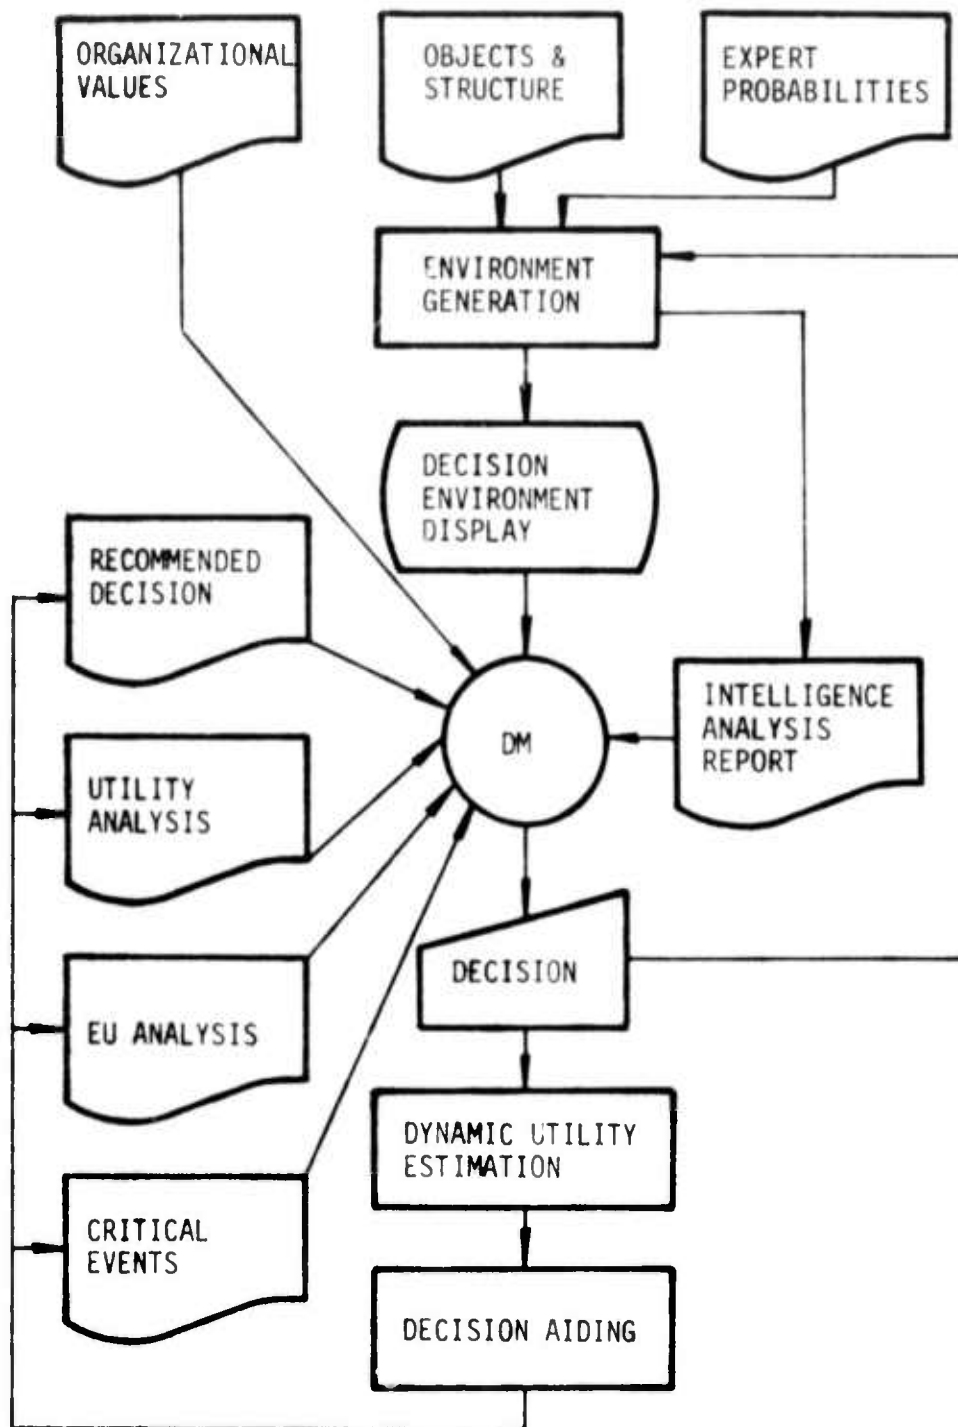
FIGURE 3-5. ADDAM FUNCTIONAL ORGANIZATION

utilities in the discriminant function (i.e., the pattern weights) are adjusted by means of an adaptive error-correction training algorithm. Thus, the utility estimator "tracks" the operator's decision behavior and "learns" his behavior. Unlike a conventional pattern classifier, however, the training process is not terminated once it is fully trained. This enables the utility estimator to respond dynamically to changes in the operator's patterns of behavior (i.e., his utilities). This adaptive decision model and the significance of the estimated utilities are discussed in Section 3.3.

Decision Aiding. Decision aiding is provided to the operator as an option under experimenter control. Decision aiding, in the current implementation, is limited to recommended decisions. The adaptive decision model, using the current estimates of the operator's utilities, computes the expected utilities of all decision alternatives and recommends the decisions which maximize the expected utility at each board location.

3.3   Dynamic Utility Measurement

In the present decision task simulation, the decision maker (DM) must make a series of interacting decisions. During the decision cycle, the operator evaluates the adequacy of existing data for tracking the objects. From this data he forms an information acquisition strategy for placing sensors to obtain additional information. He then integrates the newly obtained sensor data with the existing data to make a status decision. The information acquisition decision in a given cycle is dynamically influenced by the status decision in the previous cycle because the status decision effects the data available for the subsequent cycle.

The decisions modeled by ADDAM are decisions to place one of nine types of sensors (including a null sensor) at each of the twenty-five sectors in the simulated ocean map. A decision to place a trawler sensor, for example, in a given sector, can have one of two possible outcomes: "positive", indicating the presence of a trawler in the sector, and "negative", indicating the absence of a trawler. Since sensors are not

perfectly reliable their response may be erroneous. The decision to accept or reject a sensor outcome, however, is a component of the status reporting decision and only indirectly affects the sensor deployment decision. Accordingly, the utility estimates are a measure of the relative worth to the operator of the information outcomes of each sensor placement. Thus the utilities associated with a trawler sensor are for "positive" and "negative" intelligence about the presence of a trawler at locations where the sensor is deployed.

The utility estimates reflect the subjective values of the information outcomes in the simulation context. These values are influenced by the known cost of obtaining the information, by the perceived value of the information outcomes taking into account observed fleet behavior (element separation, patterns of movement, etc.), by the perceived reliability of the sensor (since true reliability is factored out in the utility calculation itself), and by the payoff or penalty for correctly or incorrectly reporting the status of the fleet. Other subjective factors may also be involved.

Utility Estimation. Because of the complexity of utility assessment techniques, most applications of decision theory to real world problems involve a two step process. The first step is to assess the DM's utilities and the second is to apply them to the decision problem. Since it is not feasible to re-assess utilities frequently in dynamic tasks, it is assumed that they remain static during application. Such an assumption might be valid for a static decision task. However, there is no reason to assume that the DM's utilities remain static during the performance of a multistage decision task where the environment, or the DM's perceptions of that environment are changing. For this reason, the ADDAM system uses a dynamic utility estimation technique based on a trainable multi-category pattern classifier. As the DM performs the decision task, the on-line utility estimator observes the operator's choices among the R possible decision options available to him

and views his decision making as a process of classifying patterns of event probabilities. The utility estimator then attempts to classify the event probability patterns by means of an expected utility evaluation or discriminant function.

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:

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

$$+ (1-p_i(L)) \, (1-p_{\beta k}) \, {}_kU_i^-] - C_k \qquad (3-1)$$

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

k   = 1, 2, ..., R

3-14

The utility estimator classifies the event probability patterns by using a maximum expected utility rule. These classifications are compared with the operator's decision and, whenever they are incorrect, an adaptive error-correction training algorithm is used to adjust pattern weights (corresponding to utilities). In this manner, the utility estimator "tracks" the operator's decision making and "learns" his utilities.

The utility estimator transforms the DM's patterns of behavior into a matrix of utilities which characterizes this behavior. Thus, the utility matrix of an operator who adopts a "cost be damned" strategy will be different from an operator who behaves in a miserly fashion. An operator focusing on icebergs will have a pattern different than one focusing on nets.

The utility matrix (Figure 3-6) 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 affect its usage by the DM, and this will be reflected in the estimates of his utilities for information from that sensor.

Model Validity. In training the utility estimator, the utilities are adjusted relative to each other until the model is able to predict the operator's decisions. Thus the utility estimates from one sensor can be compared with those from another to analyze the relative worth of outcomes

from the two sensors. Likewise, analysis of the overall matrix can provide information about the DM's overall strategy. In analyzing the utilities, however, great care must be exercised. The utilities are derived from the DM's choice behavior through the use of a decision model and not vice versa. We cannot say that the DM behaves in a particular manner because his utilities are such and such. However, to the extent that the utilities do characterize his behavior we can use them to analyze his behavior.

One advantage of the adaptive technique for utility estimation is that a means of validation is inherent in the training algorithm. The model predicts the DM's behavior. If the predictions are correct, no training takes place; if they are incorrect, the utilities are adjusted. Thus, if the utilities converge to constant values they have perfect predictive validity. However, given the limitations of human memory, information processing, etc., it would be unreasonable to expect perfect consistency in a task as complicated as intelligence gathering. But it is reasonable to expect that as the DM learns the task and approaches a steady state behavior, the variability of the utility estimates approaches a steady state. If the operator behaves "most of the time" in a manner which is consistent with the model, the amount of variability will be small. If his behavior is erratic there may be a great deal of variability. A measure of the changes in the utility matrix, therefore, can be used to evaluate the validity of the utilities. The UMD (Utility Matrix Difference) score described in Freedy, Weisbrod, Davis, May, and Weltman (1974) is one such measure.

One way to analyze the decision maker's behavior is to assume a normative model based on the utility estimates. The model-predicted decisions are then assumed to be optimal and the difference between the expected utility of a particular model-predicted decision and the DM's actual decision is calculated. This measure has been particularly useful in evaluating the effects of decision aiding on the DM's behavior.

OBJECT TYPE

SENSOR TYPE	ICEBERG	TRAWLER	NET
NULL	--	--	--
T	--	$_T U^+_T$	--
V1	--	$_{V1} U^+_T$	$_{V1} U^+_N$
V2	--	$_{V2} U^+_T$	$_{V2} U^+_N$
N	--	--	$_N U^+_N$
I1	$_{I1} U^+$	--	--
I2	$_{I2} U^+$	--	--
S	$_S U^+_I$	$_S U^+_T$	$_S U^+_N$
E	$_E U^+_I$	$_E U^+_T$	$_E U^+_N$
NULL	--	--	--
T	--	$_T U^-_T$	--
V1	--	$_{V1} U^-_T$	$_{V1} U^-_N$
V2	--	$_{V2} U^-_T$	$_{V2} U^-_N$
N	--	--	$_N U^-_N$
I1	$_{I1} U^-_I$	--	--
I2	$_{I2} U^-_I$	--	--
S	$_S U^-_I$	$_S U^-_T$	$_S U^-_N$
E	$_E U^-_I$	$_E U^-_T$	$_E U^-_N$

U^+ SUB MATRIX

Utility for Information that an Object is Present

U^- SUB MATRIX

Utility for Information that an Object is Not Present

FIGURE 3-6.   UTILITY MATRIX

4. EXPERIMENTAL STUDY

4.1 Hypotheses

The purpose of this study was to verify the previous findings that decision aiding improves operator consistency and results in maximum expected utility decisions. In addition, data were collected to evaluate the convergence rate of the adaptive utility adjustments. An important requirement in an adaptive decision aiding system is that it respond rapidly to the operator's choice behavior and converge to a stable level that reflects this choice behavior. Specifically, the following hypotheses were tested:

1. Aided operators will be more consistent in their performance as a group, i.e., they will show less extremes of behavior than a control group.

2. Aided operators will maximize the expected utility of their decisions to a greater extent than control subjects.

3. Aiding in the form of sensor deployment recommendations will increase operator throughput per unit time.

4. The adaptive aiding system will predict a high percentage of the operator responses.

4.2 Variables and Measures

4.2.1 Treatment Groups. The independent variable in this study was the presence or absence of decision aiding in the form of recommended sensor deployment. These recommendations were based on the subject's own

4-1

utilities in a maximum expected utility model. The experimental group performed the decision task while receiving aiding in the form of recommended sensor deployment, while the control group performed the decision task without benefit of aiding.

4.2.2 Performance Data. The performance data collected from each subject during the experimental sessions included: (1) the location and type of sensors deployed, (2) the predictions made by the model, and (3) the changes in the utility values when they occurred.

4.2.3 Performance Measures. The measures of interest in the experiment are the utility values, the utility matrix difference (UMD) score, the frequency of model predictions, the decisions per unit time, and the deviation from maximum expected utility. The measures of utility values, UMD score, and frequency of model predictions, when observed over time, indicate the extent of model adaptation to the subject's decision behavior.

The utility matrix difference (UMD) score is a measure of the variability of the utilities. This measure is computed as follows:

$$UMD(t_1, t_2) = \sum_{k,i} |{_kU^+_{it_2}} - {_kU^+_{it_1}}| + \sum_{k,i} |{_kU^-_{it_2}} - {_kU^-_{it_1}}|$$

It measures the variability of the utility values from cycle t_1 to t_2.[1] In the following analysis, however, 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

---

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

start of the session, t_0, to the end of the session, t_e. It is defined as:

$$SUMD = \sum_{t=t_0}^{t_e-1} UMD(t, t+1)$$

The deviation from maximum expected utility, the primary dependent variable, is a measure of the subject's deviation from optimal decision behavior which, by definition, is that which maximizes the expected utility of his decisions. Since a major objective of decision aiding is to improve DM performance, this becomes a measure of the effectiveness of a decision aiding scheme. The mean deviation from maximum expected utility for the test session is calculated as follows:

$$DEVMAXEU = 1/N \sum_{i=1}^{N} | EU_i \text{ (recommended)} - EU_i \text{ (taken)}|$$

where N = number of recommended decisions and the recommended decision is the maximum EU decision.

Decision per unit time is a measure of the subject's throughput during the experimental session. It, too, is a measure of the effectiveness of decision aiding. Since each subject received a test session of fixed duration the number of status reports filed, i.e., the number of decision cycles completed, was used for this measure.

4.3   Subjects

Twelve male subjects were recruited from nearby Naval reserve units. The subjects ages ranged from 23 to 44 years, while their classification ranged from E3 to E8. These subjects were chosen to be representative of the potential users of computer aided decision and control devices in the military.

4.4   Experimental Design

A simple one-way experimental design was used. One randomly selected group received decision aiding during the experimental session; the other group served as control and did not receive aiding.

4.5   Procedure

Each subject received four sessions of 1-1/2 hour duration. The first three sessions were training sessions. The training included instructions on system operation, actual handling of the equipment to familiarize the subjects with the input formats and other task features, and general strategy instructions. The training procedure was the same for all subjects. The subjects were instructed to deploy sensors according to the general strategy described in Table 4-1. The subjects were informed that part of their task was to decide the correspondence between the values in the intelligence report and the categories of the strategy.

For the fourth session, subjects were randomly divided into control and experimental groups. Those in the experimental group were given an indoctrination explaining how sensor recommendations (aiding) were actually controlled by the subject's own behavior, as well as instruction on the importance of the final session. The control group was given instruction only as to the importance of the final session.

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

TABLE 4-1. GENERAL DECISION STRATEGY GUIDELINES

Probability of Iceberg	Sensor Type
Low	e
Med	i1
High	i2

Probability of Trawler	Sensor Type
Low	t
High	s

Probability of Net	Sensor Type
Low	n
High	n

Probability Both Trawler and Net	Sensor Type
t Higher	v1
n Higher	v2

*Use the e sensor if there is the possibility that the iceberg
will be in the same location as the trawler/net.

5. EXPERIMENTAL RESULTS

5.1 General Observations

The subjects learned the task procedure readily and by the middle of the second training session could efficiently handle most of the task situations. The system responded to the subject choices rapidly as each new situation appeared and by the end of the third training session was able to predict the individual subject's choice behavior in a great percentage of the situations that arose.

5.2 Decision Aiding

5.2.1 Effect on Variability. The performance of the experimental subjects who received aiding, and the control subjects who did not is shown in Table 5-1. The tabulated score is the mean deviation from maximum expected utility (DMEU). As a group, the aided subjects were markedly more similar in the decision-making performance, while the control group showed more extremes of behavior. An F-test comparing the variance of the two groups was significant at the .001 level.

5.2.2 Effect on Deviation. Table 5-1 shows that on the average, the aided operator performed much closer to maximum expected utility, i.e., he deviated less from his own optimum norm. Indeed, the significant inter-group variability of the control group was due to some members producing gross deviations from optimum performance. Because of the heterogenity of variance a log transformation of the DMEU scores was performed. A subsequent F-test on a log transformation between the group means was significant at the .05 level.

5.2.3 <u>Effect on Decision Speed</u>. Table 5-2 shows the number of completed task cycles (status reports filed) in the fourth session. The aided group had a significantly greater mean output F-test, p < .025) and less variable performance.

5.3 <u>System Adaptive Characteristics</u>

In addition to subject performance data, the characteristics of the utility estimation program were evaluated. An important requirement in an adaptive decision aiding system is that it responds rapidly to the operator's decision-making preference and converges to a stable level that reflects these preferences. The lack of significant lag in adaptive estimation is an important requirement in a dynamic situation in which the operator changes his approach to meet new contingencies. Once the utility estimation program has acquired the decision maker's preferences it should converge to stable values and provide acceptable aiding recommendations.

Figure 5-1 shows the mean utility adjustments per decision as a function of decision cycles for the trawler sensor (this sensor was deployed 3 to 4 times per cycle). Initially each allocation results in a utility adjustment. Eventually, adjustment is made only about once in ten allocations. There appear to be two stages in machine adaptation: (1) a rapid stage, in which the major portion of adaptation is made; and (2) a gradual stage, in which minor adjustments are accomplished. In this typical case, major adaptation was completed in only five decision cycles.

The magnitude of the utility adjustments for the t trawler sensor for one subject is shown in Figure 5-2. The utility of true negative trawler sensor responses shows the greater magnitude of adjustment since the operator deployed the trawler sensor only in low probability occurrences.

TABLE 5-1
PERFORMANCE CONSISTENCY
MEAN DEVIATION FROM MAXIMUM EU

	NO AIDING (CONTROL)	AIDING
OPERATOR 1	33	65
2	1013	12
3	306	18
4	607	1
5	127	48
6	5397	109
STANDARD DEVIATION	2064	40.5
MEAN	1247	42.

TABLE 5-2
NUMBER OF STATUS REPORTS FILED

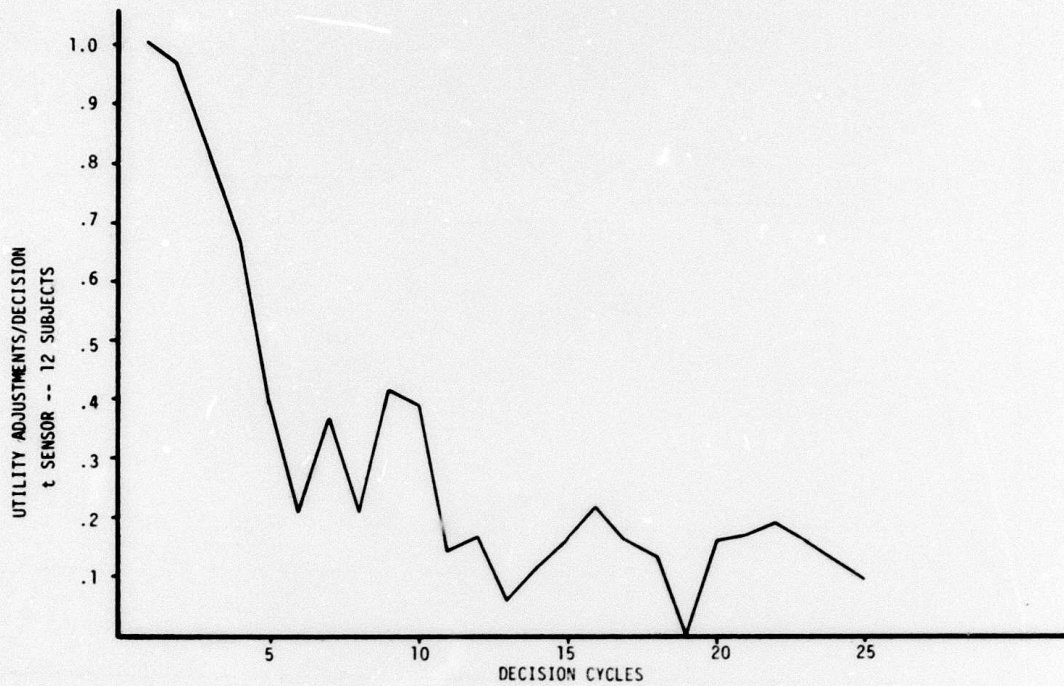| OPERATOR	NO AIDING	AIDING
1	47	52
2	38	59
3	66	66
4	43	64
5	35	59
6	53	63
MEAN	47	61
VARIANCE	128	25
VARIANCE/MEAN	2.7	0.4

FIGURE 5-1
UTILITY ADJUSTMENT/DECISION FOR THE DECISION CYCLES ON WHICH
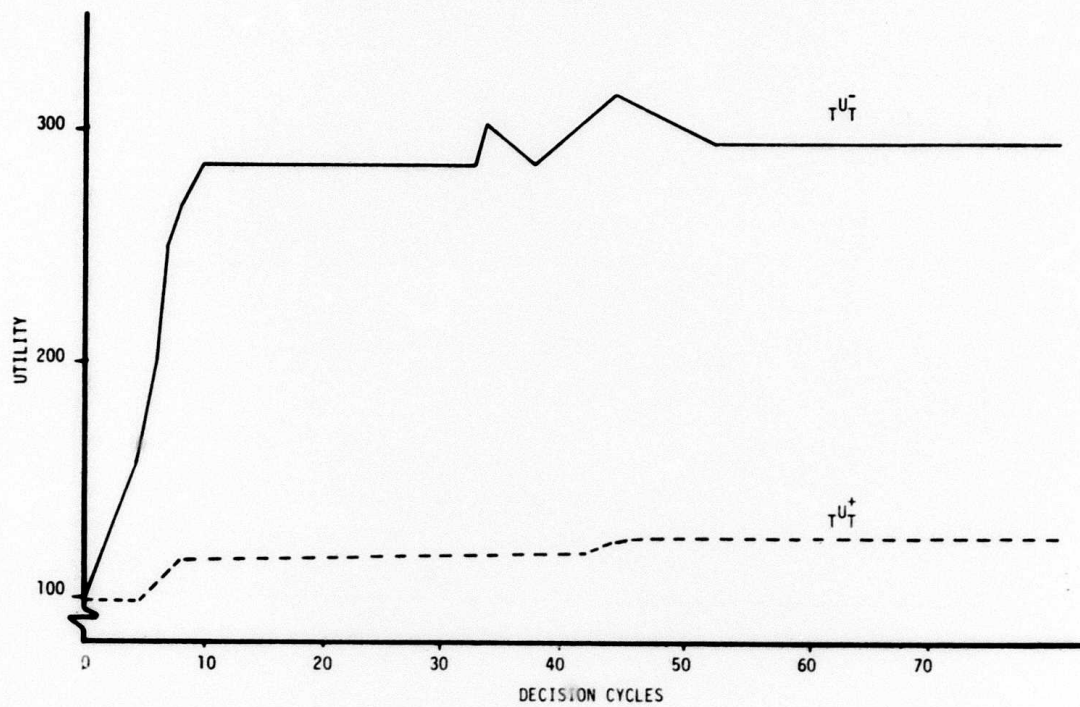THE TRAWLER SENSOR UTILITY WERE ADJUSTED (12 SUBJECTS)

FIGURE 5-2
MAGNITUDE OF TRAWLER SENSOR UTILITY ADJUSTMENT
AS A FUNCTION OF DECISION CYCLES (SUBJECT: PC)

These data show the utilities of this subject were rapidly adjusted until the tenth decision cycle after which they remain at stable values.

A more global measure of the system is the UMD score which is the magnitude of the adjustment of all utilities for each of the 8 sensor types. Figure 5-3 shows the UMD score based on a ten trial interval for one subject's three training sessions.

The magnitude of the utility adjustments for all the sensors being used initially decreases over the first 30 decision cycles as the various types of situations occur in the scenario and the subject's associated preferences are assessed. Subsequently the UMD score converges to relatively stable values as small adjustments are made to handle singularities in the scenario and associated decision preferences.
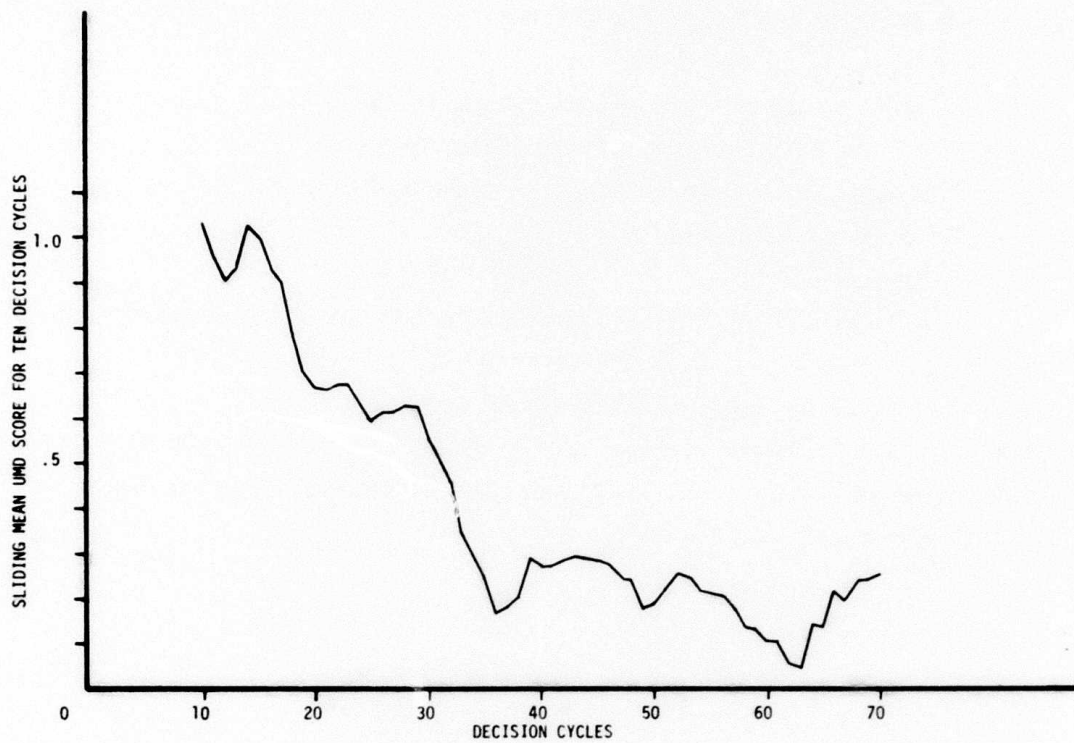
FIGURE 5-3
UMD SCORE AS A FUNCTION OF DECISION CYCLE
FOR SUBJECT HH

5-8

6. CONCLUSIONS AND RECOMMENDATIONS

6.1 Dynamic Utility Estimation

The significance of the adaptive aiding approach using an expected utility model of the operator lies in its appropriateness, systemicity and testability. In the present decision task operators readily adopted a strategy which involved preferences of resource allocation as related to objectively determined probabilities. The decision behavior was observed concurrently by an on-line utility estimation program. The utilities converged to stable values rapidly and the EU model of the operator was capable of predicting the greater majority of the operator's decisions. The advantages of this technique compared to the more formal techniques have been discussed above.

6.2 Requirements for Adaptive Aiding

To provide adequate decision aiding in a dynamic environment an adaptive system must be capable of recognizing and responding to important changes in operator strategy that are necessitated by sudden contingencies in the environmental situation or due to evolving organizational constraints and stabilizing at values which accurately reflect the operator's preferences. The ADDAM system is capable of meeting the general adaptive and accuracy requirements of real world systems to which it would be applied.

6.3 Aiding and Decision Performance

Adaptive aiding does improve the operator's decision making performance in terms of consistency. In many decision-making situations researchers have observed that subjects alter their approach using inappropriate criteria. They tend to change their decision making tactics

due to the outcomes of the preceding four or five trials, which is usually an inadequate sample in a stochastic environment. This bias in judgements of representativeness is found in subjects even when the role of sample size is formally emphasized (D. Kahneman and A. Tversky, 1972).

The belief in the representativeness of small samples has been found in technically educated and experienced subjects in military decision making studies (R. C. Sidorsky and S. R. Simoneau, 1970) and in professional researchers (A. Tversky and D. Kahneman, 1971). The point is that adaptive decision aiding ameliorates these biases and allows the operator to obtain a more representative sample to test the usefulness of the strategy he has developed.

The decision aiding also improves the operator's decision output. Part of the improvement is the time saved from manual formatting of the input. Another contributing factor is the greater relative efficiency of recognition as a memory process where aiding is present as compared to total recall required without aiding.

Finally, decision aiding can minimize the effect of habituation and of errors of omission on decision making performance, since the aiding mechanism has different limits of memory and attention than the human operator. This trait of completeness is an important factor in effective decision making performance.

6.4   Application

Adaptive aiding, like other forms of decision aiding, is applicable to a wide variety of systems in which the deficiencies of human decision makers may be overcome by techniques to augment human memory and logic

processes. Adaptive aiding is particularly applicable in situations where it is observed that humans often solve complex problems that do not yield to analysis or to strict adherence to doctrine or standard operating procedures. An organization which imposes values and goals in terms of general guidelines but otherwise allows the individual some autonomy in developing a viable strategy is the type of system that would benefit most from the adaptive aiding approach.

One of the advantages of an adaptive aiding system is a cognitive one. Studies of man-machine interaction with computer decision aids have indicated that acceptance of the aid is dependent on whether they agreed with the logic or process used by the aiding device or if they have a choice in the logic used (R. A. Hanes and J. W. Gebhard, 1966). An adaptive system, used by operators with knowledge of its nature, results in a high degree of apparent control and acceptance of the aiding system.

6.5 Future Research Direction

Future studies will utilize the payoff score and cost feedback feature recently implemented to investigate the effectiveness of aiding under constraints of limited resources and organizationally imposed values. The work as of this date has primarily focused upon demonstrating the effectiveness of aiding in reducing the subject's deviation from maximum EU decision making, and the basic validity of the model in predicting DM choices with a directly imposed strategy framework. Future work will be carried out within the context of limited resources with the organizational values imposed indirectly through a payoff matrix involving cost. Thus, the DM is free to evolve his approach and a successful approach may be more parsimonious and give more license to divergence than a maximum expected utility model would strictly allow. The goal of future work is to obtain data on the limits of the EU model. This work should suggest heuristic approaches that could improve the model's ability to predict and aid the DM in difficult decision situations.

7. REFERENCES

1. Beach, B.H. Direct and Indirect Methods for Measuring Utility. Scientific Interim Report, Department of Psychology, University of Washington, ONR N00014-67-A-0103-0011/NR 151-313, Personnel and Training Research Programs Office, Office of Naval Research, July 4, 1975.

2. Brown, R.V., Hoblitzell, C.M., Peterson, C.R., and Ulvila, J.W. Decision and Designs, Inc., Final Technical Report No. 1, No. 74-2, N00014-74-C-0263, Code 455, Office of Naval Research, September, 1974.

3. Dawes, R.M. Graduate Admission; A Case Study. Oregon Research Institute, Technical Report 10(1), 1970.

4. Edwards, W. Dynamic Decision Theory and Probabilistic Information Processing. Human Factors, 1962, 4:59-74.

5. Fishburn, P.C. Methods of Estimating Additive Utilities. Management Science, March, 1967, Vol. 13 (7): 435-53.

6. Freedy, A., Weisbrod, R., Davis, K., May, D., and Weltman, G. Adaptive Computer Aiding in Dynamic Decision Processes: Adaptive Decision Models and Dynamic Utility Estimation -- Part I of a Two-Part Report. Perceptronics Technical Report. No. PTR-1016-74-5(I). N00014-73-C-0286/NR196-128, ARPA Order No. 2347, Code 455, Office of Naval Research, May 1, 1974.

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

8. Howard, Ronald A. The Foundations of Decision Analysis. IEEE Transactions on Systems Science and Cybernetics, September, 1968. Vol. SSC-4 (3): 211-219.

9. Kahneman, D., and Tversky, A. Subjective Probabilities: A Judgment of Representations. Cognitive Psychology, 1972, 3:430.

10. Kelly, C.W. and Peterson, C.R. Probability Estimates and Probabilistic Procedures in Current Intelligence Analysis, Report FSG 5047, Federal Systems Division, IBM Corp., January, 1971.

11. Kneppreth, N.P., Gustafson, D.H., and Leifer, R.P. Techniques for the Assessment of Worth. Technical Paper 254, Con. No. DAHC19-72-6-0026, U.S. Army Research Institute for the Behavioral and Social Sciences, August 1974.

12. Margolis, M. and Leondes, C.P. On the Theory of Adaptive Control Systems: The Learning Model Approach. Proc. 1st International Congress of International Federation of Automatic Control, Moscow, USSR, 1960.

13. Payne, J.R. Miller, A.C., and Rowney, J.V. The Naval Task Force Decision Environment. Stanford Research Institute, Technical Report NWRC-TR-8, SRI Project No. 3315, N00014-74-C-0357/NR 274-199, Office of Naval Research, September 1974.

14. Sidorsky, R.C., and Simoneau, G. An Experimental Evaluation of TACTRAIN: An Approach to Computer-Aided Tactical Decision-Making Training. Decision-Making Study, Final Report 1. Naval Training Device Center (Orlando, Florida), NAV-TRADEVCEN Report 1329-4, March, 1970.

15. Tversky, A. Additivity, Utility, and Subjective Probability. Journal of Mathematical Psychology, 1967, 4, 175-202.

16. Tversky, A. and Kahneman, D. Belief in the Law of Small Numbers. Psychological Bulletin, 1971, 76:105.

17. Weisbrod, R. Davis, K., Freedy, A., and Weltman, G. Adaptive Computer Aiding in Dynamic Decision Processes. An Initial Study in Dynamic Utility Convergence and Decision Aiding. Perceptronics Technical Report No. PTR-1016-74-11, N00014-73-C-02861 NR 196-128, ARPA Order No. 2347, Code 455, Office of Naval Research, November 29, 1974.