

**Best Available
Copy
for all Pictures**

AD-780 953

ADAPTIVE COMPUTER AIDING IN DYNAMIC
DECISION PROCESSES. PART I. ADAPTIVE
DECISION MODELS AND DYNAMIC UTILITY
ESTIMATION

Amos Freedy, et al

Perceptronics, Incorporated

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

1 May 1974

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

AD-780 953

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER PTR-1016-74-5(1)	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES: Part I - Adaptive Decision Models and Dynamic Utility Estimation		5. TYPE OF REPORT & PERIOD COVERED Semi-Annual Technical Report 10/1/73 - 4/1/74
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Amos Freedy, Richard Weisbrod, Kent Davis, Donald May, Gershon Weltman		8. CONTRACT OR GRANT NUMBER(s) N00014-73-C-0286
9. PERFORMING ORGANIZATION NAME AND ADDRESS Perceptronics, Inc. 6271 Varie! Avenue Woodland Hills, Ca. 91364		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Work Unit ARPA Order No. NR 196-128 2347
11. CONTROLLING OFFICE NAME AND ADDRESS		12. REPORT DATE May 1, 1974
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		13. NUMBER OF PAGES 64
		15. SECURITY CLASS. (of this report) Unclassified
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release; Distribution Unlimited		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None		
Reproduced by NATIONAL TECHNICAL INFORMATION SERVICE U S Department of Commerce Springfield VA 22151		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) COMMAND CONTROL COMPUTER AIDED DECISIONS ADAPTIVE LEARNING SYSTEMS MAN-MACHINE INTERACTION MAN-COMPUTER SYSTEMS UTILITY ESTIMATION		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes the implementation of a system for adaptive computer aiding in dynamic decision processes and provides theoretical background for some of the underlying techniques. The report is presented in two parts under separate covers. This is Part I. Part I, "Adaptive Decision Models and Dynamic Utility Estimation," includes (1) a description of the adaptive decision model for the decision		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

task; (2) a presentation of the concept of dynamic utility and a technique, based on machine learning principles, for adaptive on-line estimation of these utilities; (3) a description of the overall system and software; and (4) the overall objectives and approach for an experimental program involving the system.

Part II, "Scenario Generation by Elicited Expert Probabilities," describes a unique technique for generating realistic dynamic decision environments. It covers the following topics: (1) a review of requirements for and techniques of generating scenarios; (2) a description of the elicited expert probability approach; (3) a description of the system for generating fishing fleet scenarios and the associated hardware and software.

This report was prepared as part of a long term program directed toward the application of adaptive computer techniques for aiding the decision maker and toward the experimental investigation of the factors which influence optimal decision aiding in complex, realistic, intelligence-gathering tasks.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

PTR-1016-74-5(I)
May 1, 1974

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES:
ADAPTIVE DECISION MODELS AND DYNAMIC UTILITY ESTIMATION
(PART I OF A TWO-PART REPORT)

AMOS FREEDY, RICHARD WEISBROD
KENT DAVIS, DONALD MAY, 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. N00014-73-C-0286 (NR-196-128)*

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

PERCEPTRONICS

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

TABLE OF CONTENTS

	<u>Page</u>
List of Figures	iii
List of Tables	iv
1. SUMMARY	1-1
1.1 Purpose	1-1
1.2 Problem and Methodology	1-1
1.3 Accomplishments	1-3
1.4 Future Work	1-4
1.5 Report Organization	1-6
2. INTRODUCTION	2-1
2.1 ADDAM Overview	2-1
2.2 Organization of Part I of the Report	2-4
3. ADAPTIVE DECISION MODELS	3-1
3.1 Overview	3-1
3.2 Adaptive Expected Utility Model	3-1
3.3 Application of the Model	3-3
3.4 Estimation of Model Parameters	3-6
4. UTILITY ASSESSMENT	4-1
4.1 Techniques of Utility Assessment	4-1
4.2 Dynamic Utility Estimation	4-4
4.3 Utility Validation and Convergence	4-10
5. ADDAM: A SYSTEM FOR MAN/COMPUTER DECISION RESEARCH	5-1
5.1 Decision Task	5-1
5.2 System Hardware	5-5
5.3 Program Structure	5-5
5.4 Man/Computer Interfaces	5-10
5.5 Decision Aiding	5-17
5.6 Current Status of ADDAM	5-19
5.7 Operational Demonstration of Utility Estimation	5-19

	<u>Page</u>
6. RELEVANT ISSUES IN SYSTEM EVALUATION	6-1
6.1 Objectives	6-1
6.2 Validation of the Model	6-1
6.3 Factors Affecting Utility Estimates and Convergence	6-2
6.4 Man/Computer Performance	6-3
6.5 Initial Experimentation	6-5
7. Bibliography	7-1

LIST OF FIGURES

<u>Figure</u>		<u>Page</u>
2-1	ADDAM FUNCTIONAL ORGANIZATION	2-2
4-1	SCHEMATIC REPRESENTATION OF DYNAMIC UTILITY ESTIMATOR	4-8
5-1	DECISION TASK SEQUENCE	5-4
5-2	ADDAM SYSTEM HARDWARE CONFIGURATION	5-6
5-3	ADDAM FUNCTIONAL FLOWCHART	5-7
5-4	OPERATOR INFORMATION DISPLAY	5-11
5-5	SENSOR PLACEMENTS	5-13
5-6	INTELLIGENCE ANALYSIS REPORT	5-15
5-7	OPERATOR UTILITY REPORT	5-18
5-8	ESTIMATED UTILITIES	5-21

LIST OF TABLES

<u>Table</u>	<u>Page</u>
5-1 Sensor Properties	5-3
5-2 An Operational Experiment Decision Strategy	5-22

1. SUMMARY

1.1 Purpose

This report presents the results of the second six months of a research program directed towards the application of adaptive learning systems to aiding in dynamic decision processes. The research goals of this program are as follows:

- a. Establish a mathematical model and a system structure for on-line adaptive computer modeling and aiding in dynamic decision making.
- b. Experimentally determine the factors which influence optimal decision aiding in complex, realistic task situations.
- c. Move toward full automation of routine multivariate, judgmental decision making.

1.2 Problem and Methodology

In dealing with real world problems, decision makers (DM) must frequently respond to dynamic input environments of multivariate data. These data come from sources of differing reliabilities and costs and have different values in the achievement of decision objectives. Decisions are made sequentially, and their consequences are likely to affect future choices. The ability of the operator to develop a satisfactory strategy for relating the poorly defined inputs to his successive decisions is a major determinate of success. Learning may be a significant part of this process, particularly in non-stationary decision environments.

Military examples of such situations range from the global to the highly specific. They include DM responses to broad range and regional intelligence reports, to local command and control needs (such as

deposition of air, sea, and ground forces), to photo image interpretation, and to noisy signals characteristic of sonar and radar returns. Numerous examples occur outside of the military as well. Besides national intelligence, these include crime prevention, air and highway traffic control, population and environmental planning, and that quintessential decision problem -- the stock market.

The approach to dynamic decision making under development at Perceptronics involves the concept of a trainable parallel decision maker model which continuously "tracks" the DM's decision responses in real time, learns his decision strategy, and aids or automates the decision process as the situation requires. In effect, the experienced decision maker "shows" the computer how to optimize in his own terms. The machine then continues the process and, in turn, aids the DM or performs his role autonomously.

The ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System represents an application of this concept.

The purpose of the ADDAM System is to provide a flexible vehicle for research in areas of dynamic decision theory, adaptive decision models, dynamic utility estimation, and man/computer decision making. The system combines the following elements:

- o Dynamic Decision Environment Generator
- o Simulated Intelligence Analysis Report
- o Decision Environment Display
- o Adaptive Decision Model
- o Dynamic Utility Estimator
- o Decision Aiding Based on Utility Feedback
- o Minicomputer Implementation

The decision task of the operator is to deploy sensors of varying object sensitivity, reliability, and cost to obtain intelligence information about the behavior of a simulated fishing fleet. The task sequence consists of deploying sensors, receiving sensor outputs, reporting fleet status, receiving an intelligence report (probabilities of movements, based upon the status report), receiving aiding information (at this time, limited to sensor deployment suggestions), and again deploying sensors.

The ADDAM System has been implemented on an Interdata, Model 70 mini-computer with 24K bytes of core memory. A teletype and an IDIgraf graphic display terminal with 2K bytes of internal memory and direct memory access are used to provide a man/machine interface.

1.3 Accomplishments

The following is a summary of the accomplishments to date.

Dynamic Decision Environment Generator. A Scenario Generator has been developed to generate fishing fleet scenarios for the dynamic decision task. This generator and the routines which simulate the behavior of the sensors are operational and have been used to generate scenarios. A generalized methodology for scenario generation based upon elicited expert probabilities has evolved out of the considerable insight gained through operational testing.

Simulated Intelligence Analysis Report. A technique for using the expert's probability matrix from the Scenario Generator to estimate the probabilities of events in the real world has been developed. The probabilities are based on the status of the fishing fleet as reported by the operator, and they are presented to the operator in the form of an "intelligence analysis report." They are also used by the adaptive decision model.

Adaptive Decision Model and Dynamic Utility Estimator. The adaptive decision (expected utility) model of operator decision behavior and the dynamic utility estimator have been implemented and are now running. Operational experimentation has shown that the model is able to track utilities based upon a simulated operator decision strategy. The parameters of the utility estimator are currently being adjusted to improve its performance.

Decision Aiding and Man/Computer Interface. Decision aiding currently consists of suggesting maximum EU decisions to the operator. Other forms of aiding are under investigation. Subsystems for handling the interchange of information between the human operator and the computer are now operational. They allow the operator to deploy sensors and report status. They also display sensor outputs, status, and suggested sensor deployments on the IDIgraf graphics display terminal and print intelligence reports on the teletype.

Experimental Program. Work has begun on the first phase of the experimental program: the validation of the dynamic utility assessment technique. Operational experiments are being conducted to validate the algorithm for dynamic utility estimation and to determine the response of the model to different kinds of simulated operator strategies. This will provide experience and insights which will be needed when systematic experimentation is initiated with naive subjects. A convergence measure (a measure of validity of the model and the utility estimates) has been developed and an analysis of task related variables has begun.

1.4 Future Work

The primary research objectives of this three year study of adaptive computer aiding in dynamic decision making include the following:

Establish guidelines for the application of adaptive decision systems on the basis of mathematical considerations and related research.

Implement the most promising techniques as interactive computer programs for realistic decision making.

Explore in experimentally controlled environments the factors which influence effective monitoring, aiding, and automation by the adaptive learning programs. Formulate human factors criteria, and identify areas of program refinements.

Validate major findings by similar data acquisition in real world, "open" decision making situations.

Develop equipment design specifications (including major trade-offs) for practical field implementation of recommended techniques.

The research plan for the coming year includes the following:

Perform an experimental study which will demonstrate overall system operation in adaptive acquisition of decision strategies. estimation of operator utility, and ability to predict operator behavior.

Define a meaningful measure of convergence of subjective operator values in order to be able to validate and utilize estimated utilities.

Validate the accuracy of the EU model as a basis for estimation of operator utilities and for aiding.

Identify conditions and constraints under which model is effective.

Establish a theoretical framework for adjusting and modifying the model to most effectively predict operator decision behavior in performing complex intelligence gathering tasks.

Explore the possibilities of including factors such as operator biases and cognitive constraints in the model.

Develop and experimentally evaluate decision and aiding schemes which are based on model-derived "dynamic" utility estimates.

Establish the scope of applicability and develop guidelines for using the system in operator decision aiding and decision theory research.

1.5 Report Organization

The report is divided into two parts which are published under separate covers. The first part, Adaptive Decision Models and Dynamic Utility Estimation, presents the philosophical and theoretical basis for two unique features of the Perceptronics approach: an adaptive expected utility model and a technique for real-time estimation of dynamic utilities. It also describes ADDAM (Adaptive Dynamic Decision Aiding Machine), a system which applies these ideas to an intelligence gathering task, and outlines a plan for experimentation with the system.

The second part of the report, Scenario Generation by Elicited Expert Probabilities, describes a unique technique developed at Perceptronics for generating realistic dynamic decision environments and the application of the technique to generating the intelligence gathering task.

2. INTRODUCTION

2.1 ADDAM Overview

The ADDAM (Adaptive Dynamic Decision Aiding Mechanism) System is a flexible vehicle for conducting research in areas of dynamic decision theory, adaptive decision models, dynamic utility estimation, and man/computer decision making.

The relationship among the basic elements of ADDAM are shown schematically in Figure 2-1. A dynamic decision environment is probabilistically generated on the basis of expert probabilities and an organization structure specified by the experimenter. This decision environment is displayed to the decision maker as seen through costly, unreliable sensors which he has deployed to gather intelligence information. On the basis of this sensor information, an intelligence analysis report, and varying forms of decision aiding, the operator makes decisions to deploy new sensors and to report the status of the environment. Finally, the operator's decision behavior is analyzed using pattern classification techniques to dynamically estimate his utilities for intelligence information. These utilities form the basis for several forms of decision aiding.

Dynamic Decision Environment Generator. Scenarios of events in a dynamic decision environment are generated by a unique application of Bayesian information processing techniques. Unlike PIP systems (Edwards, 1962), which aggregate conditional probabilities elicited from experts in order to estimate the probabilities of complex events in a real world, the aggregated probabilities are used to obtain a Monte Carlo simulation of the real world. Scenarios thus generated have statistical consistency and can appear to respond dynamically to decision outcomes. The environment generator is used to generate scenarios involving the movements of a fishing fleet. Only those features of the scenario which are detected by the operator's sensors are actually displayed.




Figure 2-1. ADDAM FUNCTIONAL ORGANIZATION

Simulated Intelligence Analysis Report. The statistical consistency of the environment generator permits the simulation of a 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 more 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.

Dynamic Utility Estimator. The dynamic utility estimator, in conjunction with an adaptive decision (EU) model, employs the principle of a trainable multi-category pattern classifier to assess the operator's utilities. The utility estimator uses the expected utility model as an evaluation function for classifying patterns of event probabilities into decision categories. The utilities are adjusted adaptively by means of an error correction training procedure which makes the classifications more descriptive of the decision maker's (DM's) behavior. Since this training is done continuously as the task is being performed, the system is able to "track" changes in the DM's utilities in real time.

Decision Aiding Based on Utility Feedback. ADDAM provides a framework for forms of decision aiding and training based upon the concept of utility feedback. Utility feedback is made possible by the availability of real-time utility assessments. Forms of aiding include recommending optimal decisions, analyzing DM utilities for information, analyzing his expected utilities for information, and highlighting critical events. Training can come from comparisons of operator utilities with expert utilities or organizational values. Both decision aiding and training are adapted to the individual because they are based upon his personal value structure.

2.2 Organization of Part I of the Report

The ADDAM System employs several unique concepts, including that of adaptive decision models and dynamic utility estimation. The adaptive decision model concept is examined in Chapter 3 and dynamic utility estimation is the main focus of Chapter 4. Chapter 5 describes the implementation of the ADDAM System and briefly discusses the results of a preliminary operational experiment. Chapter 6 outlines objectives for planned experimentation with ADDAM.

The system for generating the dynamic decision environment is described in Part II of this report (under separate cover).

3. ADAPTIVE DECISION MODEL

3.1 Overview

The ADDAM System uses an expected utility (EU) model as a basis for estimating utilities and aiding a decision maker in the performance of a dynamic decision task. The EU model is unique in one important aspect: its utilities are adaptively adjusted, in response to decision maker (DM) behavior during the performance of the task, by using trainable pattern-classifying system techniques. Thus, the adaptive expected utility (AEU) model, continuously tracks the operator's decision strategy as it changes in response to environmental changes, acquired learning, and other factors.

This chapter provides a conceptual framework and rationale for the use of an adaptive expected utility model. It also describes the particular AEU model used by ADDAM. Discussion of the technique for dynamic utility estimation is reserved for Chapter 4.

3.2 Adaptive Expected Utility Model

Decision models are often classified as being either descriptive or normative. Descriptive models attempt to describe the decision behavior of decision makers and predict their actions while normative models attempt to prescribe the decisions they should make in order to satisfy specific decision criteria. Thus, a decision model used for decision aiding would usually be classified as normative. However, an adaptive decision model is not so easily classified.

In tracking the operator's behavior, the adaptive decision model is acting as a descriptive model. The error-correction procedure (see Chapter 4) adjusts the model's parameters (its utilities, in the case of an AEU model) in a manner which will make the model more descriptive of the operator's decision behavior. However, when the model is used as a basis for decision aiding it is normative.

The use of an adaptive model for decision aiding establishes a complex, poorly understood, feedback-loop between the model and the human operator. The behavior of the decision model is modified by the human operator's behavior and this model behavior, in turn, is used to influence the operator's action. Establishing the nature of this "symbiotic" relationship and the important factors influencing it (see Chapter 6) is a major long term goal of our research.

Expected utility models are widely accepted as normative models for decision-making under risk (Luce and Raiffa, 1957, and Krantz, Luce, Suppes, and Tversky, 1971). The work of Tversky (1967), Goodman, Saltzman, Edwards, and Krantz (1971), and others has indicated that expected utility models provide a good first approximation to decision making under risk, at least for simple gambling situations where the number of attributes is low and the DM can relate to all attributes in terms of probabilities. Several researchers, however, have raised doubts about the model. Lichtenstein and Slovic (1971) argue that descriptive models of choice must take cognitive factors into account, and Tversky and Kahneman (1973) have shown that DMs use heuristics, termed representations, to relate the cues associated with making decisions. Wendt (1973) questions the validity of the general concept of maximization of expectation as a normative model.

Another factor to consider is the validity of using objective probabilities to describe decision behavior. The alternative of using subjective probabilities and a Subjective Expected Utility (SEU) model introduces a great deal of complexity because one must now contend with two sets of unknown variables: subjective probabilities and utilities. If it can be assumed that the subjective and objective probabilities are equal, as was done (for example) by Seghers, Fryback and Goodman (1973), then the SEU and EU models are equivalent. This assumption is reasonable if the decision maker is told the objective probabilities.

From the discussion above, it is clear that expected utility models have shortcomings. Nevertheless EU models are useful in situations where these shortcomings are not significant. The ADDAM System uses EU to provide a structure for utility assessment and decision aiding. The adaptive mechanism then searches for subjective values which predict operator behavior in terms of the EU model. It is not necessary for the EU model to be perfectly predictive of DM behavior since the adaptive mechanism responds to patterns of behavior. Individual DM actions tend to average out. Inconsistencies between the predictions of the model and the behavior of the DM will cause the value of the estimated utilities to fluctuate over time, but as long as this variance stays within reasonable bounds the model utilities can be used as a relative measure of the actual DM utilities (see Chapter 4).

3.3 Application of the Model

The adaptive expected utility model is applied to modeling DM behavior in the performance of an intelligence gathering task. Briefly (see Section 5.1 for a more complete description), the task consists of deploying sensors with differing object sensitivities, reliabilities, and costs to gather intelligence information about a simulated fishing fleet environment.

The expected utility model is based upon the utility for information from each kind of sensor. The model is first expressed for the most general case where both alpha (false negative) and beta (false positive) errors are possible. It is then simplified to the form used in ADDAM, in which only beta errors can occur. In its most general form, 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:

$$\begin{aligned}
 EU_k(L) = \sum_i k^M_i [& p_i(L)(1-k^p_\alpha) k^U_i && \text{(True Positive)} \\
 & +(1-p_i(L))(1-k^p_\beta) k^U_i && \text{(True Negative)} \\
 & -(1-p_i(L)) k^p_\beta k^{U_{ie}} && \text{(False Positive)} \\
 & -(1-p_i(L)) k^p_\alpha k^{U_{ie}} && \text{(False Negative)} \\
 & -C_k && \text{(Sensor Cost) (2-1)}
 \end{aligned}$$

The symbols used above are defined as follows:

$EU_k(L)$ = expected utility of deploying sensor of type k at location L

k^M_i = sensor capability mask bit. 1 if sensor is capable of reporting information about attribute i , 0 if sensor is incapable of reporting about attribute i .

$p_i(L)$ = probability of an object with attribute i at location L .

k^p_α = probability of an α error (false negative) from a sensor of type k .

k^p_β = probability of a β error (false positive) from a sensor of type k .

k^U_i = utility of correct information about attribute i from a sensor of type k .

$k^{U_{ie}}$ = utility of erroneous information about attribute i from a sensor of type k .

C_k = fixed cost of deploying a sensor of type k .

In ADDAM's task simulation the assumption is made that the sensors are extremely sensitive and never fail to detect the presence of an object, i.e., $k p_{\alpha} = 0$. However, because of this sensitivity, they often detect objects which are not present, i.e., $k p_{\beta} \neq 0$. This assumption results in the following simplifications to the EU model:

$$EU_k(L) = \sum_i k^M_i [p_i(L) k^{U_i} + (1-p_i(L))(1-k p_{\beta}) k^{U_i} - (1-p_i(L)) k p_{\beta} k^{U_{ie}}] - C_k \quad (3-2)$$

$$= \sum_i k^M_i \left[[1-k p_{\beta} (1-p_i(L))] k^{U_i} - [k p_{\beta} (1-p_i(L))] k^{U_{ie}} \right] - C_k \quad (3-3)$$

We let

$$k^{p_i}(L) = k^M_i [1-k p_{\beta} (1-p_i(L))] \quad (3-4)$$

and

$$k^{p_{ie}}(L) = k^M_i [k p_{\beta} (1-p_i(L))] \quad (3-5)$$

then

$$EU_k(L) = \sum_i [k^{p_i}(L) k^{U_i} - k^{p_{ie}}(L) k^{U_{ie}}] - C_k \quad (3-6)$$

The model selects, for each location on the board, the sensor whose expected utility is maximum. To prevent the deployment of a sensor at every location, a null sensor whose mask bits are all zero is included as one of the alternatives. Essentially, the cost of this null sensor acts as a threshold EU, below which no sensor is deployed.

3.4 Estimation of Model Parameters

The adaptive expected utility model makes use of three sets of parameters: probabilities, utilities, and costs. The probabilities used by the model are objectively determined values. These values are displayed to the operator. The probabilities of false alarms, $k p_{\beta}$, are characteristics of the sensors and their values are set by the experimenter. The probabilities associated with objects, $p_i(L)$, are computed on the basis of the operator's status report by aggregating the "elicited expert probability matrix." The procedure is described in Part II (Section 3.3) of this report (under separate cover).

The costs of deploying sensors are set by the experimenter during program initialization. The values may vary according to the needs of the experiment.

The utilities are the only values which are actually estimated. The values are dynamically estimated by tracking the operator's behavior as he performs the decision task. It is, in fact, the mechanism for dynamic utility estimation which makes the EU model adaptive. Utility estimation, in general, and the mechanism for dynamic utility estimation, in particular, are the topics of Chapter 4.

4. UTILITY ASSESSMENT

4.1 Techniques of Utility Assessment

The use of expected utility decision models (as well as other utility-based models) depends on being able to estimate the utilities of the decision maker. The usual procedure for applying these models to complex decisions in real world contexts involves two steps. The first step is to estimate the utilities using one of several conventional utility assessment techniques and the second step is to use these utility estimates in applying the model.

Techniques currently used for utility assessment can be divided into four categories: ordinal scale methods, direct methods, gambling methods, and multivariate methods. These techniques have been reviewed and analyzed by Kneppreth, Gustafson, Johnson, and Leifer (1973). With ordinal assessment methods, the decision maker is asked to qualitatively rank his preferences. His rankings are used to develop an ordinal scale of utilities. This can be converted to an interval scale if equal intervals are assumed, but the resulting scale is only approximate.

Direct methods of utility assessment (e.g., Beach, 1972) require the DM to make quantitative estimates of his subjective feelings. These methods are quick and easy to use since they do not require large numbers of repetitious judgments and calculations, but their validity has been questioned because they do not follow the axioms of utility theory. However, several researchers (Beach, 1972; Fisher, 1972) have shown that direct utility estimates are comparable with axiomatically derived estimates.

Gambling methods require the a priori decomposition of complex decisions into many simple lotteries. Either the probability or the outcome of each lottery is varied until the DM is indifferent between the lottery and a "sure thing". Utilities thus calculated are axiomatically valid, but the process is long, tedious, and somewhat contrived.

Multivariate methods are used to obtain utility functions which involve more than one attribute, especially when the attributes are not independent. The procedure involves determining which combinations of attributes result in indifference on the part of the DM when compared with a "reference" combination. By making a large number of such comparisons, a set of indifference curves can be developed. Making these comparisons is a long and tedious process.

Validation. 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). Because of the difficulty in obtaining independent measures of these subjective quantities, it is difficult to validate utility estimates.

One widely used method of validating utility estimates is to check for consistency. If a DM makes choices which are inconsistent with the axioms of utility theory or other requirements of the utility assessment process, the inconsistencies are called to his attention and resolved. Human decision makers, however, are not perfectly consistent (Edwards, 1961), thus it may be unreasonable to require perfect consistency in many real world decision tasks. Further, the consistency check only insures that the utility estimates are internally consistent. It does not insure that they accurately reflect the DM's true subjective values.

Comparing the operator's utilities with organizational utilities is a second method of validating utility assessments (Peterson, 1971). Organizational utilities are values which are defined by the organization to which the DM belongs. These externally specified values are then used as a standard for evaluating the DM's utility estimates. Other externally defined decision criteria can also be used to provide "objective" standards for evaluating utility estimates.

Another method of validation involves the examination of the reliability of value judgments over time. Value judgments made at one time,

which systematically differ from judgments made under the same conditions at another time would tend to invalidate both sets of value measures (Miller, Kaplan, and Edwards, 1967). Similarly consistency in behavior over time would tend to validate the value measures.

Construct validity is another means of utility validation. This is based on the idea that two different methods of measuring the same abstract quantity should give comparable results (Miller, Kaplan, and Edwards, 1967). Fischer (1972) describes a number of different comparisons which have been used. They include the degree of correlation (Fischer uses the word convergence, but we will reserve this term for us in a more mathematical sense) between (1) wholistic (intuitive) judgments and those based upon decomposition techniques (of utility assessment); (2) model predicted choices and real choices; (3) values obtained using two or more different utility assessments; and (4) different subjects.

In light of the definition of utility as a measure which characterizes a person's judgments, it would seem that the degree of correlation with actual behavior would provide the strongest kind of validation. However behavioral validation must be used with caution. Characterizations of human decision behavior (i.e., utility assessments) must be made within the context of a model of that behavior (in most cases some sort of EU model is used). Thus, the validity of utility estimates is inherently limited by the validity of the decision model. Behavioral validation, more than other methods of validation, calls attention to the limitations of the model.

Static vs. Dynamic Utilities. Because of the complexity of utility assessment techniques, most applications of decision theory to real world problems involves 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. Because it is not feasible to re-assess utilities frequently in repetitive tasks, it is assumed that they remain static during this 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. Nor is it reasonable to assume that they remain the same when the context changes from that of a set of lotteries to the real world task.

In performing a multistage (dynamic) decision task the DM acquires information which affects his subsequent performance. Bayesian information processing systems (Edwards, 1962) make use of this information to modify the probabilities associated with the decision processes. This information may change the decision maker's goals as well. Thus, the relative values he assigns to decision outcomes will change. Also, changes in the probabilities of events may have an effect on the utilities of alternatives, contrary to the usual assumption of independence between probability and utility (Slovic, 1966). Thus, dynamic decision tasks introduce a need for dynamic utility assessment techniques.

The following section introduces an adaptive technique for dynamic utility assessment which was developed at Perceptronics.

4.2 Dynamic 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 R possible decision options available to him, viewing 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. These classifications are compared with the operator's decisions and an adaptive error-correction training algorithm is used to adjust pattern weights, which correspond to utilities, whenever the classifications are incorrect. Thus, the utility estimator "tracks" the operator's decision making and "learns" his utilities.

Pattern classification techniques have been used in a limited fashion to perform decision making functions. For example, Henderson (1972) used a two-category classifier for diagnostic evaluation of medical questionnaires and Bartels and Wied (1974) used a multi-category classifier for evaluation of microphotometric measurements in clinical cytodiagnosis. In both of these typical cases the classifiers respond to pattern cues which have an objectively "correct" classification. These correct responses were learned during an off-line training period and then applied to the performance of a static decision task.

The application of pattern classification techniques to utility estimation was suggested by Siagle (1971), who pointed out that the utility function was an evaluation function and that the function could be learned from a person's preferences. A two-category pattern classifier which adaptively estimates operator utilities for computer or human control of a man/computer decision task was developed by Freedy, Weisbrod, and Weltman (1973). It was shown in pilot studies that this utility estimator could track the operator's utilities on-line during on-the-job performance of the decision task and that the estimator was dynamically responsive to changes in the operator's value structure (Weltman, Steeb, Freedy, Smith, and Weisbrod, 1973). In this case, the pattern classifier responded to patterns of probabilities (of human and computer success and failure) and the operator's subjective preferences for human or computer control of the task.

Multi-Category Pattern Classifiers. A multi-category pattern classifier (Nilsson, 1965) receives patterns of data and responds with a decision to classify each of the patterns in one of R categories. The classification is made on the basis of R linear discriminant (or evaluation) functions, each of which corresponds to one of the R categories. The discriminant functions are of the form

$$g_i(\bar{X}) = \bar{W}_i \cdot X \quad \text{for } i = 1, 2, \dots, R \quad (4-1)$$

where $\bar{X}$ is the pattern vector and $\bar{W}_i$ is a weight vector. The pattern classifier computes the value of each discriminant function and selects the category, i , such that

$$g_i(\bar{X}) > g_j(\bar{X}) \quad (4-2)$$

for all $j = 1, 2, \dots, R; i \neq j$.

The adaptive error-correction training algorithm is very straightforward. Whenever the category selected by the pattern classifier, i , is different from the actual classification, k , the weights $\bar{W}_i$ are adjusted to reduce (punish) the value of $g_i(\bar{X})$ and the weights $\bar{W}_k$ are adjusted to increase (reward) the value of $g_k(\bar{X})$. Thus,

$$\bar{W}_i' = \bar{W}_i + d \cdot \bar{X} \quad (\text{Reward}) \quad (4-3)$$

$$\bar{W}_k' = \bar{W}_k - d \cdot \bar{X} \quad (\text{Punish}) \quad (4-4)$$

where d is the correction increment.

A linear pattern classifier is trained by presenting it with a "training set" of preclassified patterns. These patterns are presented to the machine, one at a time, until it is able to classify them perfectly. Once the machine is trained, it is then used to classify patterns which have not previously been classified. If the categories are linearly separable the training procedure is guaranteed to find a set of solution weight vectors in a finite number of steps (Nilsson, 1965) and this solution set will yield a zero error rate. If the categories are not linearly separable, the error rate will not be zero, though it may be satisfactorily low (Slagle, 1971), and training will have to be terminated after some finite number of steps.

The Dynamic Utility Estimator. The dynamic utility estimator, shown schematically in Figure 4-1, classifies pattern vectors

$$\bar{P} = [p_{1,1}, p_{2,1}, \dots p_{i,k} \dots] \quad (4-5)$$

whose components, $p_{i,k}$, are the aggregated probabilities of the i th decision outcome, as influenced by the reliability of the k th sensor. These components correspond to the probabilities of correct and incorrect sensor responses defined in Equations 3-4 and 3-5.

The discriminant functions are the expected utilities of each sensor decision as defined in Equation 3-6. The utility estimator computes the EU of each sensor at each location on the board and selects those sensors (including the null sensor described on page 3-5) for which the EU is maximum. The selected sensor at each location is compared with the actual decision 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.

The training procedure for the utility estimator is as follows. Whenever the decision, j , selected by the utility estimator differs from the decision, k , selected by the operator, the utilities associated with the estimator decision are punished and those associated with the operator decision are rewarded:

$$\begin{aligned} jU_i^{t+1} &= jU_i^t - d \cdot jP_i(L) \\ jU_{ie}^{t+1} &= jU_{ie}^t + d \cdot jP_{ie}(L) \quad (\text{Punish}) \end{aligned} \quad (4-6)$$

$$\begin{aligned} kU_i^{t+1} &= kU_i^t + d \cdot kP_i(L) \\ kU_{ie}^{t+1} &= kU_{ie}^t - d \cdot kP_{ie}(L) \quad (\text{Reward}) \end{aligned} \quad (4-7)$$




Figure 4-1. SCHEMATIC REPRESENTATION OF DYNAMIC UTILITY ESTIMATOR

The utilities at time $t+1$ are computed for all decision outcomes (attributes sensed), i . The correction increment, d , is a constant which can be adjusted to give optimum convergence of the utility estimator. The other values are as defined in Section 3.3 (page 3-5).

The procedure for training the utility estimator differs from the procedure for a pattern classifier. The primary output of a pattern classifier is the classification of the pattern. The pattern weights are usually of no importance by themselves. Thus, a pattern classifier is trained only until it is able to make classifications with an acceptable degree of accuracy. With the utility estimator, on the other hand, the pattern weights (utilities) are the primary outputs and the classifications are of secondary importance since they are used only for training and some forms of decision aiding. Thus, the utility estimator is trained continuously so that it can track the operator's utilities as they change in response to the dynamics of the task.

Because the utility estimator is being continuously trained it would be useful to examine the behavior of the utilities under various conditions. If the probability patterns are linearly separable into categories (decisions), the utility estimator will learn to classify them perfectly after a finite number of steps. Since training takes place only when there are classification errors, each utility will converge to a single value. If the operator's values change, the utility estimator will begin making errors again, training will take place, and the utilities will converge to a new set of values.

If the patterns are not linearly separable, a different situation arises since the utility estimation can never learn to classify perfectly. In a conventional pattern classifier, linear inseparability is reflected in the error rate. In a system which is continuously being trained, this error rate keeps the utilities from converging to a single value. However, the utilities may approach a steady state value within a range of variance.

4.3 Utility Validation and Convergence

The primary means of validating the dynamic utility estimates would be to demonstrate that they characterize the operator's behavior within the context of the decision model. This type of validation is inherent in the error-correction procedure of the dynamic utility estimation technique.

The predictive validity of the utility estimates is a matter of degree. Perfect predictive validity would require that the operator's behavior in the task be perfectly consistent with the decision model. Perfect predictive validity would result in the perfect convergence of the utilities. Given the limitations of human memory, information processing, etc., it would be unreasonable to expect this in a task as complicated as intelligence gathering. Thus, the primary demonstration would be to show that, as the operator 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.

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

$$UMD(t_1, t_2) = \sum_{k,i} |kU_i^{t_2} - kU_i^{t_1}| + \sum_{k,i} |kU_{ie}^{t_2} - kU_{ie}^{t_1}| \quad (4-8)$$

The UMD is a global measure of the variance of the utility values from time t_1 to time t_2 . The magnitude of the UMD provides a measure of the validity of the EU model and of the utilities. The rate of change of the UMD indicates the stability of the utility estimator. As the utility estimator approaches a steady state, the rate of change of the UMD will approach zero.

5. ADDAM: A SYSTEM FOR MAN/COMPUTER DECISION RESEARCH

The purpose of the ADDAM (Adaptive Dynamic Decision Aiding Machine) System is to provide 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 with an adaptive decision model, a system for dynamic utility estimation, and mechanisms for man/computer interaction and decision making.

5.1 Decision Task

The initial decision task simulation is a simplification of the intelligence gathering task described by Freedy, Weisbrod, May, Schwartz, and Weltman (1973). The task involves deploying sensors of varying object specificity, reliability, and cost in order to gather intelligence information about a dynamically varying hierarchical organization -- a fishing fleet. In performing this task, the operator (decision maker) must report what he believes to be the status of the task environment.

The Environment. The environment is a homogeneous expanse of ocean. This expanse, referred to as the board, is divided into a five by five square two-dimensional spatial grid. The fishing fleet, consisting of trawlers which may or may not deploy nets, moves around the board, from square to square. Also present are icebergs which similarly move around the board. These objects are constrained to the board.

There are several environmental conditions which affect the behavior of the objects. These include time of day (day or night), weather conditions (clear or stormy), and phase of moon. The presence of nearby icebergs also effects the behavior of trawlers.

Each object on the board has the following characteristics associated with it: object type (iceberg, trawler, trawler with nets deployed), location, and heading (North, East, South, West, Null). These objects cannot be seen by the operator except through sensors which he has deployed.

The Sensors. The properties of the sensors available to the operator include object sensitivity, response specificity, error rate, and cost. Object sensitivity refers to the kinds of objects which the sensor can detect. Response specificity refers to the sensor's ability to identify the objects which it has detected. Error rate, at the present time, is limited to false positive (β -error) rates. It is also possible to specify a false negative (α -error) rate; however, the decision model is currently implemented to include only β -errors.

These properties permit the specification of a wide variety of different kinds of sensors. Table 5-1 defines the set of sensors used for initial testing of the system. This set includes two trawler sensors with different false alarm rates and costs, a net sensor, an iceberg sensor, and an "everything" sensor. All of these sensors respond with the kind of object detected. On the other hand, the "something" sensor, a low cost, low reliability sensor, is sensitive to all kinds of objects, but only responds positively or negatively.

The operator has an unlimited number of sensors of each type at his disposal. However, he can deploy only one sensor per square, and he must pay a cost for each sensor he deploys. The sensor only responds to objects within its square.

The Decision Task Sequence. The decision task sequence (Figure 5-1) begins when the operator deploys his sensors. Once he has finished deploying his sensors, he receives a report of the sensor outputs. Some of these sensors may give a positive response while others may not. On the basis of the sensor responses, knowledge of sensor behavior, previous

Table 5-1
Sensor Properties

<u>Type</u>	<u>Object Sensitivity</u>	<u>Response Specificity</u>	<u>False Alarm Rate</u>	<u>Cost</u>
T1	Trawler	Trawler	0.10	2.50
T2	Trawler	Trawler	0.30	1.50
N	Trawler with Net	Net	0.30	1.50
I	Iceberg	Iceberg	0.20	2.00
E(Everything)	Trawler/Net/Iceberg	Trawler/Net/Iceberg	0.01	5.00
S(Something)	Trawler/Net/Iceberg	Positive/Negative	0.40	0.50




Figure 5-1. DECISION TASK SEQUENCE

sensor responses, etc., the operator reports what he believes is the status of the environment. This status report, which includes object type, location, and heading, is used by the system to generate an intelligence analysis report.

The intelligence analysis report gives the probabilities that each square will contain an object on the next turn. For simplicity, only squares with non-zero probabilities are reported. This report is displayed to the operator and is used by the adaptive EU model. The operator then receives aiding information which will help him make his next set of sensor deployment decisions. Finally, he deploys sensors to begin the cycle anew.

5.2 System Hardware

The ADDAM System is implemented on an Interdata Model 70 minicomputer with 24K bytes of core memory. The man/computer interface is through a teletype and an Information Displays, Inc. IDIgraf graphic display terminal with 2K bytes of internal memory and direct memory access. Figure 5-2 illustrates the physical arrangement.

The hardware was selected to provide the capability for real time operation of the system. The operator inputs his decisions, and receives sensor outputs, intelligence reports, and aiding information within a short period of time. The main time limitation is the speed of the teletype in printing out intelligence reports. A hardware "precision interval clock" is used to control the amount of time allocated for input of sensor and status decisions, and for experimental sessions.

5.3 Program Structure

The program is organized as a set of functional modules which are controlled by a Master System Scheduler. The functional flow chart (Figure 5-3) illustrates the sequencing of the basic modules. Additional modules (not illustrated) can be used to perform such functions as computing the



Figure 5-2. ADDAM SYSTEM HARDWARE CONFIGURATION




Figure 5-3. ADDAM FUNCTIONAL FLOWCHART

status board payoff or the convergence measures, or gathering statistical data for evaluating experimental results.

Master System Scheduler. This program consists of a control structure for scheduling a sequence of calls to functional module subroutines. It sets up the communications between subroutines and then passes control to them. The Master System Scheduler also contains a clock routine for allocating time for the performance of certain functions. For example, if the experimenter wishes to give the operator one minute to input all of his sensor or status decisions, the clock routine will terminate the inputs after one minute. If the experimenter wants the whole session to last 15 minutes, the clock routine handles it.

Initialize Program. This module sets up the program for the beginning of a run. It initializes the clock, and allows the experimenter to input the expert probabilities used to generate the environment and the initial starting values for the utilities. If the operator has used the system before, it is possible to input his previous utilities as a starting point.

Update Environment State. This module, the Scenario Generator, generates the decision task scenario, one step at a time, from the matrix of elicited expert probabilities. The technique used to generate the scenario is described in Part II of this report. This module also contains the sensor routines which act as windows through which the operator can observe the state of the environment.

Display Sensor Outputs. This routine displays the sensor outputs on the IDIgraf graphics terminal. Nothing is displayed during the first pass through the Master System Scheduler control loop before the first sensors have been deployed. The display formats for all display routines are described in Section 5.4.

Receive and Display Status Decisions. This routine sets up a communications link with the operator in order to receive status decisions. The operator types in a status decision on the IDIgraf keyboard and transmits it to the computer. The decision is then displayed in the status input area of the display screen. When the operator manually terminates the input, the status decisions are displayed on the situation board (on the IDIgraf screen).

Generate Intelligence Data. This module analyzes the operator's status report and generates the probabilities used in the intelligence report. These probabilities are generated using the current status of the environment, as reported by the operator, and the Elicited Expert Probability Matrix used by the scenario generator.

Display Intelligence Report. This module arranges the intelligence data into report format and prints it out on the teletype.

Compute Suggested Sensor Decisions. This module is the heart of the Adaptive Expected Utility Model. It computes the expected utility of using each sensor and selects the sensors which maximize the EU at each board location. One of the sensor choices is a null sensor which does nothing and is not reported (displayed) to the operator. The cost of deploying the null sensor acts as a threshold EU, below which no sensor is deployed.

Display Suggested Sensor Decisions. This routine displays the suggested sensor decisions on the IDIgraf terminal. It does not display decisions to deploy null sensors. This routine can be disabled by the experimenter for experiments where it is not desirable to display the suggested sensors.

Receive and Display Sensor Decisions. This module functions in the same manner as the Receive and Display Status Decisions module.

Train Utilities. This module is the heart of the Utility Learning Machine. It compares the sensor decision suggested by the adaptive EU model with the decision made by the operator and rewards or punishes the appropriate utilities.

5.4 Man/Computer Interfaces

Human interaction with the ADDAM System takes place on two levels. The first level involves the operator (experimental subject) interface with the system during the performance of the decision task. As far as the operator is concerned, this interface has a fixed structure during the performance of the task. The system requires certain kinds of inputs from the operator and it, in turn, provides him with specific kinds of outputs.

The second interface level is between the experimenter and the system. At this level, the experimenter is allowed a great deal of flexibility in modifying the nature and complexity of the task environment, the performance characteristics of the decision model, and structure of the operator/computer interface.

The first part of this section will describe the current structure of the operator/computer interface. The second part will describe the degree of flexibility available to the experimenter.

Operator/Computer Interface. The operator interaction with the system begins with a request for sensor decisions. The system displays the "Sensor Deployment" heading on the IDIgraf graphics display terminal and positions the cursor to the start of the first input line. The operator then inputs the location (square coordinates) and type for the first sensor to be deployed and transmits it to the computer. The system responds by positioning the cursor to the start of the next input line and the process is repeated. An example of the inputs are shown in the upper right hand corner of the display shown in Figure 5-4.




Figure 5-4. OPERATOR INFORMATION DISPLAY

If the system has suggested sensor decisions, the input procedure is a bit different. The recommended actions are printed on the teletype as shown in Figure 5-5. Also, the suggested sensor decisions are displayed in the sensor input area of the display and the cursor is positioned in front of the first decision. The operator can input Y or N to accept or reject the suggestion. If he inputs C (change), the suggestion is erased and the cursor is positioned to input the operator's changes. When the operator reaches the bottom of the list he can input additional sensor decisions. Thus, the operator processes the suggested decisions in checklist fashion, accepting, rejecting, or changing them, and adding new decisions to the bottom of the list.

Once the operator has made all of his sensor decisions, he terminates the input and the sensors appear in the upper left hand corner of the selected squares on the board. If a sensor detects an object, it begins to blink and the sensor output appears to the right of the entry in the sensor deployment list (see Figure 5-4).

Following sensor deployment, the system will accept the operator's status decisions. The "Fleet Status" heading appears on the display and the cursor appears at the start of the input line. The operator types in the location, object type, and heading of the first object to be indicated and transmits it to the computer. The system responds by positioning the cursor to the start of the next input line and the process is repeated. Sample inputs are shown in the lower right corner of Figure 5-4. When the operator terminates his inputs, the status indicators appear on the board and the listing of status decisions disappears from the display. The sensor decision listing also disappears, but the sensors themselves do not disappear from the board.

The intelligence analysis report, printed on the teletype, includes the location and the probability of occurrence for each type of object.

RECOMMENDED SENSOR PLACEMENT IS

FIELD SECTOR	SENSOR TYPE
A2	E
A3	T1

Figure 5-5. SENSOR PLACEMENTS

Locations for which the probabilities are zero are omitted from the report. A sample intelligence report is illustrated in Figure 5-6.

Experimenter/Computer Interface. The experimenter interacts with the ADDAM System primarily to adjust the system's behavior to meet the needs of his experiments. The system is designed to allow the experimenter to modify, with a minimum of effort, the nature and complexity of the decision task environment, the decision model performance characteristics, and the structure of the operator/computer interface.

The experimenter controls the decision task environment by modifying the characteristics of the scenario and the environment sensors. The scenario is generated from an object list and a matrix of conditional probabilities of transformations which determine the behavior of these objects (see Freedy, May, Weisbrod, Weltman, 1974). The experimenter can modify the behavior of the objects by changing the conditional probability values. He can add new objects, for example, additional trawlers or a factory ship, by making additional entries in the tables which define the object list.

Control over the properties of sensors is given to the experimenter in two ways. Object sensitivity and response specificity are defined by tables in the sensor routines. These properties can be changed by modifying the table and new sensors can be defined by adding new entries to the table. False alarm rate and sensor costs are specified by the experimenter during the initialization of the program. Thus they can be reset at the start of each run.

The performance characteristics of the decision model can be controlled by modifying (1) the initial utility values used by the adaptive EU model, (2) the learning rate of the utility estimator, and (3) the EU evaluation function used by both the model and the utility estimator. The easiest to modify are the initial values of the utility matrix, which are input by

FISHING FLEET INTELLIGENCE REPORT

REPORT NO. 24

ESTIMATED ELEMENT DISTRIBUTION FOR NEXT PERIOD IS

FIELD SECTOR	PROBABILITY THAT SECTOR OCCUPIED BY			
	ICEBERG	TRAWLER	TRAWLER AND NET	SOME ELEMENT
A1	8	0	0	8
B1	45	0	0	45
A2	45	0	0	45
B4	0	0	2	2
A5	0	0	3	3
B5	0	80	10	82
C5	0	0	3	3

Figure 5-6. INTELLIGENCE ANALYSIS REPORT

the experimenter during program initialization. These initial values affect the behavior of the adaptive EU model and the utility estimator, at least during the early stages of a run. Since the initial performance of a decision aiding system can have a significant influence on the operator's use of the decision aid (Halpin, Thornberry, and Streufert, 1973), the choice of initial utility values may be very important in some experiments. The choice of initial utilities might be made from a standard set of values (e.g., all values equal to one), a set of values learned during a previous run with the same subject, or a set of "expert" utilities.

The learning rate of the utility estimator is controlled by the correction increment, d , defined in Chapter 3. This parameter affects the rate of convergence of the utility estimator and determines its sensitivity to changes in the operator's decision behavior. The value of the correction increment also affects the amount of variance which will result from inconsistent operator behavior.

Modification of the expected utility function (Equation 3-6) is the most difficult to use of the three methods of controlling the decision model. This expression, which is also used as a discriminant function by the utility estimator, is programmed into the system. Reprogramming is facilitated, however, by the modular design of the system. Such redefinition of the EU function might be done if, for example, new kinds of objects (e.g., factory ships) were introduced into the scenario generator.

Changing the structure of the operator/computer interface is accomplished primarily by changes to the Master System Scheduler. Since the scheduler is a sequence of calls to functional modules, inserting and deleting calls to modules will change the structure of the interface as seen by the operator. For example, in the initial experiments to validate the model, it is not desirable to aid the operator by suggesting sensor decisions to him. This change is easily implemented by deactivating the routine which displays the suggested decisions. Other changes, such as displaying or not displaying the operator's payoff score, are similarly accomplished.

The amount of time available to the operator is an important part of his interaction with the system. Parameters in the Master Scheduler determine how much time is allowed for the operator to make sensor decisions and status decisions. Other parameters determine whether the input periods are terminated by operator action, expiration of the time period, or both. Also determined by parameters in the scheduler is the amount of time allocated for an experimental run.

5.5 Decision Aiding

One aspect of the man/computer interaction which is of central importance is decision aiding. One type of decision aiding, the suggesting of sensor decisions on the basis of maximum expected utility, has been implemented and is currently being investigated. A number of other forms of decision aiding are also of interest. An analysis of the operator's immediate value structure is one such form of aiding. This permits both self and outside assessment of the operator's decision behavior, as well as comparison with other value standards such as organizational values or expert opinion. Figure 5-7 illustrates a utility report which is now displayed to the experimenter at the end of a run. This report could also be presented to the operator as a rudimentary form of aiding. Changes in the dynamic estimates of the operator's utilities and inconsistencies in his behavior may signal significant happenings in the decision environment or a major reassessment by the operator of important decision criteria. These changes could be reflected in a similar report.

Analysis of the expected utilities of information and decisions is another form of aiding which may be of value to the operator. Highlighting "important" incoming data is a form of decision aiding which may prevent the operator from missing critical events. The decision to continue to acquire information or to report it may depend upon the instantaneous EU of the information and, perhaps, some threshold of importance.

OPERATOR UTILITY REPORT

PERIOD NO. 23

SENSOR TYPE	POS UTILITY CORRECT INFO			NEG UTILITY FALSE ALARM		
	ICEBERG	TRAWLER	TRAWLER AND NET	ICEBERG	TRAWLER	TRAWLER AND NET
I	226	0	0	95	0	0
T1	0	170	0	0	98	0
T2	0	100	0	0	100	0
N	0	0	147	0	0	97
E	136	140	139	95	99	98

Figure 5-7. OPERATOR UTILITY REPORT

Finally, sensor decisions which are optimal on the basis of criteria other than maximum expected utility can be suggested to the operator. These criteria could include published policy, expert consensus, or objective performance measures, as well as operator utilities. Continued acceptance of such machine suggestions could lead to a progressive transfer of the task to the computer, with the human operator retaining the capability to review and override machine decisions.

5.6 Current Status of ADDAM

The ADDAM System has been implemented and is now running. It is currently being used for operational experiments and shakedown tests. The Scenario Generator, used to generate fishing fleet scenarios for the dynamic decision task, and the routines which simulate the sensors are operational and have been used to generate scenarios. A modification of the Scenario Generator is being used to generate intelligence reports based upon expert probabilities and the operator's status report.

The adaptive decision model and the dynamic utility estimator are operational, as is the man/computer interface subsystem. Minor revisions are being made as the nature of the interaction becomes more apparent; however, major changes are not anticipated. Decision aiding currently consists of suggesting maximum EU decisions to the operator. Other forms of aiding are under investigations.

5.7 Operational Demonstration of Utility Estimation

The adaptive decision model and dynamic utility estimator are currently being tested in operational experiments. The results of one such dynamic utility estimation test are presented in Figure 5-8. The estimated utilities for information about the fleet elements are shown as a function of the trial number. The objective of the test was to ascertain the

capability of the system to track and estimate the operator's utilities for information sources, given consistent DM behavior.

An arbitrary operator decision strategy was chosen in which decision alternatives were selected solely as a function of the probability of the movement of an object to a board location. This strategy is summarized in Table 5-2. Whenever the intelligence report showed that the probability of an iceberg at a given location to be higher than 0.60, the operator was instructed to deploy an iceberg sensor at that location. When the probability of a trawler was greater than 0.40 he was told to deploy a trawler (T1) sensor. When the probability of a trawler with a net was greater than 0.80, a net sensor was to be deployed.

At the beginning of the decision task all utilities were arbitrarily set at 100. As the decision behavior was tracked, the values separated into three distinct levels. A trend toward separation was apparent after only 15 trials. After about 40 trials the utilities converged to three distinct levels with variations around mean values.

The convergence of the utility estimates to distinct levels indicates that the EU model reflects gross operator behavior. The variations around these levels represent slight inconsistencies in the behavior. The horizontal line segments occur during periods in which no training took place. This absence of training indicates that the model predicted decisions were in agreement with the decisions made on the basis of the operator's decision strategy.




Figure 5-8. ESTIMATED UTILITIES

Table 5-2
An Operational Experiment Decision Strategy

Deploy Sensors as Shown

<u>Prob. Between</u>	<u>Iceberg</u>	<u>Trawler</u>	<u>Trawler and Net</u>
0-20	None	None	None
20-40	None	None	None
40-60	None	T1	None
60-80	I	T1	None
80-100	I	T1	N

6. RELEVANT ISSUES IN SYSTEM EVALUATION

6.1 Objectives

The objectives of the experimental programs correspond to three phases of investigation and development of the Expected Utility Decision Aiding concept.

The objectives of the three phases of investigation are:

1. Validation of dynamic utility assessment.
2. Characterization of system sensitivity to DM behavior.
3. Optimization of DM behavior through aiding.

The first objective essentially consists of demonstrating that the model is capable of predicting DM behavior with a reasonable degree of accuracy in the absence of any aiding to DM. The second objective is achieved by demonstrating the model's sensitivity to individual DM values and/or to the organizationally imposed task values. The third objective is demonstrated by showing increases in DM's consistency and performance when he is provided with various aiding information in a form that allows him to overcome some of his limitations.

The first two objectives are of immediate concern since they provide the primary validation of the approach. The emphasis during the initial phases will be investigation of the task related variables discussed below. Once the basic soundness of the approach has been demonstrated, emphasis will be shifted to the system interface variables that influence man-machine interaction with the goal of optimizing system performance.

6.2 Validation of the Model

The demonstration of the convergence of the utility matrix will validate the on-line estimated utility model. Because the adaptive utility estimates are adjusted to reflect the operator's decisions, the utilities, by definition,

correlate with those decisions. The degree of correlation is inversely proportional to the variance of the estimated utilities. Thus, a stable utility matrix reflects a high correlation with a consistent EU decision strategy. The demonstration of convergence of the utility matrix is an example of construct validity with the operator's decisions used as the comparison standard.

6.3 Factors Affecting Utility Estimates and Convergence

The factors affecting utility estimates and convergence may be divided into two categories: task related variables and system interface variables. Task related variables include:

1. Organizational or institutional values.
2. Sensor characteristics.
3. Interrelationship of environmental events.
4. Operator access to environmental information.

The organizational or institutional values relate to budget, importance of objects and events and allowable false alarm rate for the system. These, in turn, are related to the sensor characteristics of cost, object specificity, and false alarm rate. The effect of such external values are studied by using indoctrination and debriefing techniques that impose restraints and/or emphasis on various aspects of the operator's task. These institutional values act as driving functions that guide the operator's behavior to a consistent level. They differentially affect the utilities associated with specific sensor types.

The characteristics of the sensors (specificity, cost, false alarm rate) affect the degree to which a given sensor type contributes to task accomplishment in relation to other sensor types. The extent to which the degree of effectiveness among sensor types may be discriminated by the operator with given organizational values will effect the consistency of

his decisions and thus the degree of convergence of the operator model. If the characteristics of the sensors in relation to a specific task are such that the differential effectiveness among sensor types is not apparent to the operator, then his choices of sensor deployment will be less consistent. If the discriminability of sensor effectiveness is high, the operator's alternatives are less equivocal, thus facilitating decision consistency.

The greater the degree of predictive interrelationship of environmental events, the greater the potential cues on which to base decisions. Randomness of the environmental elements precludes cognitive structuring by the decision maker and contributes little toward operator consistency.

Potentially, the interrelationships of the environmental elements may provide information germane to the decision process. However, DM's access to this information is determined in the most part by the successful deployment of sensors. The extent to which DM is able to gain access to these predictive relationships is an important determinant of consistent operator behavior.

6.4 Man/Computer Performance

The third phase objectives of improving human decision making capabilities involves the investigation of variables that effect human information processing in manned systems.

These system interfaces relate to the man-machine interaction and include:

1. Aiding information
2. Attitudinal factors
3. Degree of Apparent Control

A major machine oriented variable is the type and configuration of aiding information that would be provided to improve disability in a dynamic-complex decision environment.

Information generated by the model may be used to aid the decision maker in several ways; these include:

1. Recommending sensor deployment based on the EU model of operator behavior.
2. Emphasizing important environmental events as determined from previous operator responses.
3. Calling attention to inconsistencies in value structure to prompt selective reevaluation of the action and to motivate consistent behavior.

Once the basic validity of the dynamic utility assessment approach is demonstrated, the effectiveness of the various aiding techniques will be investigated in terms of the stability of the operator's decision making behavior.

The major operator variables of interest that would affect convergence are attitudinal and situational variables that bias human interaction with machine components in a man-machine system. Consider the case in which the operator has learned to play the fishing fleet game when the only aiding being given is the intelligence report. When the aiding is introduced after the operator has reached steady state performance it will be met with a variety of reactions related to the operator's initial attitude concerning computerized systems. With an initially strong positive attitude it is expected that the operator would accept the aiding readily and the u-matrix would become stable. With a negative initial attitude the aiding would be rejected and u-matrix stability would be retarded. In the first case, with strong positive attitude, if the system performance decreased as a result of accepting aiding, the positive attitude may be replaced by a negative attitude and instability would result.

In the absence of strong initial attitudes it is expected that situational variables would influence operator behavior to a greater extent. The perceived amount of control over the decision processes utilized by

the aiding mechanism or the operator's understanding and agreement with the decision logic used is an important situational variable (Hanes and Gebhard, 1966). This variable may be manipulated in indoctrination procedures which give differential training on the nature of the adaptive process. It is expected that the perception of apparent control by the indoctrinated group will result in more frequent acceptance of aiding and concurrently a more stable u-matrix. It is also hypothesized that if system performance decreases with the acceptance of aiding, there will be a greater tendency in the indoctrinated group to continue to accept the aiding rather than reshaping the u-matrix.

6.5 Initial Experimentation

The operational experiment described in Section 5.7 provided an examination of the convergence of the utility matrix given operator behavior that is consistent with the probabilities stated in the intelligence report. However, operator behavior may not necessarily reflect such a system-internal consistency. Thus it is necessary to examine the convergence of the utility matrix given operator behavior which is consistent with imposed organizational rules and values. In this instance, operator behavior will not be entirely consistent with the intelligence report probabilities since these probabilities do not reflect false alarm sensor errors. These errors cause the operator to revise sensor deployment decisions based upon expected subsequent sensor outputs, rather than based strictly upon the intelligence report probabilities.

The current investigation includes the examination of utility matrix convergence, given operator behavior consistent with imposed rules of decision strategy. Several system-naive subjects are included in the investigation to provide a measure of system stability across decision strategies.

7. BIBLIOGRAPHY

1. Bartels, Peter H., and George L. Wied, "Scanning Microphotometry and Computer-Aided Diagnosis in Clinical Cytology," presented at a tutorial seminar on Optics in Diagnostic Medicine, University of Arizona, Tucson, January 9-11, 1974.
2. Beach, Barbara 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, 1972.
3. Edwards, Ward, "Behavioral Decision Theory," Annual Review of Psychology, 1961, 12:473-489.
4. Edwards, Ward, "Dynamic Decision Theory and Probabilistic Information Processing," Human Factors, 1962, 4:59-74.
5. Fischer, G.W., "Four Methods for Assessing Multi-Attribute Utilities: An Experimental Validation," Technical Report 037230-6-T, Engineering Psychology Laboratory, University of Michigan, ONR N00014-67-A-0181-0034/NR 197-014, Engineering Psychology Programs, Department of the Navy, September, 1972.
6. 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, ARPA, Dept. of Defense, May 1, 1974.
7. Freedy, Amos, Richard Weisbrod, Donald May, Stephen Schwartz, and Gershon Weltman, "Adaptive Computer Aiding in Dynamic Decision Processes: Analysis and Design," Technical Report PTR-73-101, Perceptronics, Inc., Woodland Hills, Calif., ONR, N00014-73-C-0286/NR 196-128, ARPA Order #2347, Dept. of Defense, October 1, 1973.
8. Freedy, A., R. Weisbrod and G. Weltman, "Self-Optimization of Task Allocation in Shared Man/Computer Control," Proceedings of the IEEE Conference on Decision and Control, San Diego, Calif., December 5-7, 1973.
9. Goodman, B., M. Saltzman, W. Edwards, and D.M. Krantz, "When SEU is Irrelevant," University of Michigan, Department of Psychology, 1971.

10. 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.
11. Henderson, Charles, "A Trainable Pattern Classifier for Medical Questionnaires," Annals of Biomedical Engineering, 1972, 1:115-133.
12. 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.
13. Krantz, D.M., R.D. Luce, P. Suppes, and A. Tversky, Foundations of Measurement: Additive and Polynomial Representations, I, New York: Academic Press, 1971.
14. Lichtenstein, S. and P. Slovic, "Reversals of Preferences Between Bids and Choices in Gambling Decisions," Journal of Experimental Psychology, 1971, 89:46-55.
15. Luce, R.D. and H. Raiffa, Games and Decisions: Introduction and Critical Survey, New York: Wiley, 1957.
16. Miller, L.W., R.J. Kaplan, and W. Edwards, "JUDGE: A Value-Judgment-Based Tactical Command System," Organizational Behavior and Human Performance, 1967, 2:329-374.
17. Nilsson, Nils J., Learning Machines, New York: McGraw Hill, 1965.
18. 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.
19. Seghers, Raymond C., Dennis G. Fryback and Barbara C. Goodman, "Relative Variance Preferences in a Choice-Among-Bets Paradigm," Technical Report No. 011311-6-T, Engineering Psychology Laboratory, University of Michigan, ONR, N00014-67-A-0181-0049/NR 197-021, ARPA Order 2105, ARPA, Department of Defense, November 1973.
20. Slagle, James R., Artificial Intelligence: The Heuristic Programming Approach, New York: McGraw Hill, 1971.

21. Slovic, Paul, "Value as a Determiner of Subjective Probability," IEEE Transactions on Human Factors in Electronics, March, 1966, HFE-7:22-82.
22. Tversky, Amos, "Additivity, Utility, and Subjective-Probability," Journal of Mathematical Psychology, 1967, 4:175-202.
23. Tversky, Amos and Daniel Kahneman, "Judgement under Uncertainty: Heuristics and Biases," Report #ORIB 13(1), Oregon Research Institute, ONR, N00014-73-C-0438/NR 197-026, ARPA Order #2449, ARPA, Department of Defense, August, 1973.
24. Weltman, Gershon, Randall Steeb, Amos Freedy, Michael Smith, and Richard Weisbrod, "Experimental Study of Man/Machine Interaction in Adaptive Computer Aided Control," Technical Report 73-10, Perceptronics, Inc., Woodland Hills, Calif., ONR Contract N00014-72-C-0093/NR 196-118, Engineering Psychology Programs, Department of the Navy, November, 1973.
25. Wendt, Dirk, "Some Criticisms of the General Models used in Decision Making Experiments," Report No. 011313-8-T, Engineering Psychology Laboratory, University of Michigan, Ann Arbor, Michigan, ONR, N00014-67-A-0181-0049/NR 197-21, ARPA Order #2105, ARPA, Department of Defense, 2 November, 1973.