AD A032796

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES:

METHODOLOGY, EVALUATION, AND APPLICATIONS

PERCEPTRONICS INCORPORATED

Amos Freedy
Kent B. Davis
Randall Steeb
Michael G. Samet
Peter C. Gardiner

DDC
RECEIVED
DEC 2 1976
A

ADVANCED DECISION TECHNOLOGY PROGRAM

CYBERNETICS TECHNOLOGY OFFICE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY
Office of Naval Research • Engineering Psychology Programs

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

BEST
AVAILABLE COPY

The objective of the Advanced Decision Technology Program is to develop and transfer to users in the Department of Defense advanced management technologies for decision making. These technologies are based upon research in the areas of decision analysis, the behavioral sciences and interactive computer graphics. The program is sponsored by the Cybernetics Technology Office of the Defense Advanced Research Projects Agency and technical progress is monitored by the Office of Naval Research – Engineering Psychology Programs. Participants in the program are:

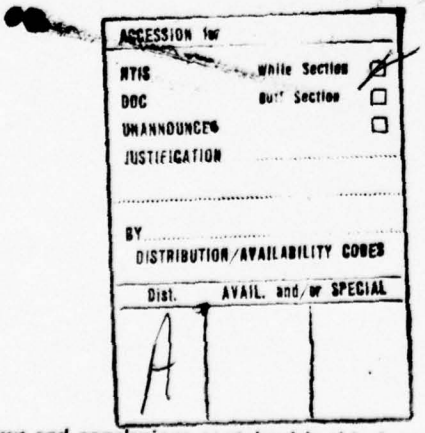
Decisions and Designs, Incorporated
The Oregon Research Institute
Perceptronics, Incorporated
Stanford Research Institute
Stanford University
The University of Southern California

Inquiries and comments with regard to this report should be addressed to:

Dr. Martin A. Tolcott
Director, Engineering Psychology Programs
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

or

Dr. Robert A. Young
Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

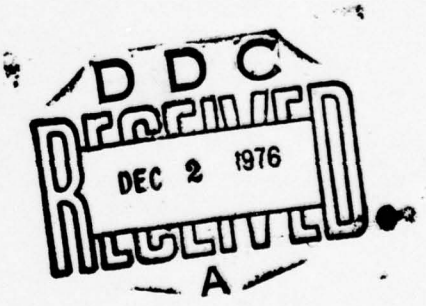


The views and conclusions contained in this document are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government. This document has been approved for public release with unlimited distribution.

FINAL TECHNICAL REPORT PFTR-1016-76-8/30
N00014-73-C-0286
NR 196-128
ARPA Order No. 2347
August 1976

ADAPTIVE COMPUTER AIDING IN DYNAMIC DECISION PROCESSES: METHODOLOGY, EVALUATION, AND APPLICATION

Amos Freedy
Kent B. Davis
Randall Steeb
Michael G. Samet
Peter C. Gardiner



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

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: PFTR-1016-76-8/30
2. GOVT ACCESSION NO.:
3. RECIPIENT'S CATALOG NUMBER:

4. TITLE (and Subtitle): Adaptive Computer Aiding in Dynamic Decision Processes: Methodology, Evaluation, and Application.

5. TYPE OF REPORT & PERIOD COVERED: Final Technical Report.

6. PERFORMING ORG. REPORT NUMBER:

7. AUTHOR(s): Amos Freedy, Kent B. Davis, Randall Steeb, Michael G. Samet Peter C. Gardiner

8. CONTRACT OR GRANT NUMBER(s): N00014-73-C-0286, ARPA Order-2347

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 NR196-128
ARPA Order #2347

11. CONTROLLING OFFICE NAME AND ADDRESS:

12. REPORT DATE: August 1976

13. NUMBER OF PAGES: 119p.

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office):
Engineering Psychology Programs, Code 455
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

15. SECURITY CLASS. (of this report): Unclassified

15a. DECLASSIFICATION DOWNGRADING SCHEDULE:

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

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

18. SUPPLEMENTARY NOTES

None

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

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

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

This is the final report of a three year program directed toward the design and development of a decision support system for tactical operations. Termed ADDAM (Adaptive Dynamic Decision Aiding Methodology), the system uses adaptive models both to capture the decision maker's strategy and to suggest actions. The report presents theoretical and methodological reviews, as well as the results of the most recent experimental study.

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

over

UNCLASSIFIED 390 285
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

LB

20. ABSTRACT (Continued)

Earlier studies of ADDAM showed that it is capable of modeling and predicting decision behavior, and of improving decision consistency. The current study extends these findings by demonstrating that adaptive aiding can also improve external performance measures in a realistic task. The experimental context was a simulated ASW task much like that faced by an ASW operations officer. The operator tracked and reported on the movements of a submarine and an interfering object, using the same types of sensors (sonobuoys, helicopters, MAD's, etc.) available in Naval ASW Exercises. The sensors varied in reliability, in specificity, and in cost. Twelve Air National Guard reservists served as subjects in the study, six without aiding and six with. Aided operators received computer assistance in the form of (1) sensor output evaluation, and (2) sensor placement recommendations. The latter were based on their own utilities for sensor information, derived in previous training sessions by means of a trainable pattern recognition program. The results showed that the aided group performed significantly better than the control group. Their total performance score improved by 88% due to gains in both decision throughput and decision quality. Also, the aided group showed greater decision consistency and lower intra-group variability.

The report discusses the practical application of adaptive decision modeling in providing a framework for decision analysis and feedback. The normative expected utility approach shows particular applicability in situations of high risk, complexity, and speed.

TABLE OF CONTENTS

		Page
1.	SUMMARY	1-1
	1.1 Purpose	1-1
	1.2 Problem	1-1
	1.3 System Concept	1-2
	1.4 Experimental Studies	1-4
	1.5 Conclusions	1-5
2.	THE PSYCHOLOGICAL BASIS OF ADAPTIVE AIDING	2-1
	2.1 General	2-1
	2.2 Man as a Suboptimal Decision Maker	2-1
	2.3 Information Gathering Behavior	2-2
	2.3.1 Description	2-2
	2.3.2 Experimental Studies	2-3
	2.3.3 Rationale for an Information Gathering Aid	2-6
	2.4 Previous Decision Aiding Systems	2-7
	2.4.1 Definition	2-7
	2.4.2 Probabilistic Information Processing	2-8
	2.4.3 Utility-Based Aiding	2-10
	2.4.4 Adaptive Decision Models	2-11
	2.5 Acceptance of Computer Aiding	2-15
	2.6 Present Adaptive Aiding Concept	2-15
	2.7 Significance of ADDAM as a Decision Aid	2-17
3.	ADAPTIVE AIDING METHODOLOGY	3-1
	3.1 ADDAM System Overview	3-1
	3.2 ADDAM Theoretical Structure	3-2
	3.2.1 Probability Aggregation	3-3
	3.2.2 Utility Estimation	3-3
	3.2.3 Strategy Recommendation	3-7
	3.2.4 Aiding in Information Acquisition	3-7
	3.2.5 Theoretical Evaluation	3-12
	3.3 Model Learning and Parameter Estimation	3-12
	3.3.1 Pattern Classification	3-12
	3.3.2 The Dynamic Utility Estimator	3-13

TABLE OF CONTENTS (Continued)

		Page
3.4	System Characteristics	3-15
3.4.1	Closed Loop Features	3-15
3.4.2	Adaptive Features	3-16
3.4.3	System Performance	3-16
3.4.4	Performance Evaluation	3-18
3.5	Decision Measures	3-20
3.5.1	Decision Outcome	3-20
3.5.2	Decision Quality	3-21
4.	EXPERIMENTAL STUDY OF ASW DECISION MAKING	4-1
4.1	Introduction	4-1
4.1.1	Overview	4-1
4.1.2	Previous Studies	4-1
4.1.3	Current Approach	4-2
4.2	Simulated ASW Task	4-3
4.2.1	The ASW Environment	4-3
4.2.2	Decision Task	4-7
4.2.3	Decision Aids	4-12
4.3	Methodology	4-14
4.3.1	Experimental Hypotheses	4-14
4.3.2	Independent Variable	4-15
4.3.3	Dependent Variables	4-15
4.3.4	Experimental Design	4-16
4.3.5	Subjects and Procedures	4-16
4.4	Results and Discussion	4-17
4.4.1	General	4-17
4.4.2	Results	4-17
4.4.3	Decision Consistency	4-19
4.4.4	Attitudes and Acceptance	4-19
4.4.5	Discussion	4-20
4.4.6	Implications	4-21
4.4.7	Value of Information	4-23
4.5	Conclusions	4-25

TABLE OF CONTENTS (Continued)

		Page
5.	TACTICAL APPLICATIONS OF ADAPTIVE DECISION AIDING	5-1
	5.1 Introduction	5-1
	5.2 The Tactical Decision Task	5-1
	5.2.1 Definition	5-1
	5.2.2 The Tactical Decision Cycle	5-2
	5.2.3 The Tactical Decision Making Environment	5-5
	5.2.4 The Tactical Decision Maker	5-5
	5.3 A Tactical Example: Anti-Submarine Warfare	5-6
	5.3.1 Task Description	5-6
	5.3.2 The ASW Decision Environment	5-8
	5.3.3 The ASW Decision Cycle	5-9
	5.3.4 Problems in ASW Decision Making	5-9
	5.4 ASW and Adaptive Decision Aiding	5-11
	5.4.1 Decision Theoretical Terminology	5-11
	5.4.2 Tactical Aiding Concept	5-13
	5.4.3 Tactical Aiding Modes	5-13
	5.4.4 Conclusions	5-16
	5.5 Illustrative Examples of Tactical Adaptive Decision Aiding	5-16
6.	TECHNOLOGY TRANSFER	6-1
7.	REFERENCES	7-1
	APPENDIX	A-1

1. SUMMARY

1.1 Purpose

This report presents the results of a three year program of research and development directed toward the design and evaluation of an adaptive decision support system for tactical operations. The system concept is based on the use of adaptive models both to capture the decision maker's strategy and to suggest actions based on this strategy. These functions are performed during actual decision making, and are specifically suited to aiding the rapid-response, dynamic decisions characteristic of tactical situations. The research goals of the program were as follows:

(a) To establish mathematical decision aiding models, operating software, and evaluation methods for on-line adaptive aiding.

(b) To determine experimentally the effectiveness and range of application of adaptive computer-based aiding in realistic dynamic task situations.

(c) To move toward operational aiding of current command and control systems using adaptive techniques.

1.2 Problem

Current tactical operations are becoming increasingly sensitive to the quality of decision making. Large stakes rest on the ability of personnel to request and process volumes of information, and to make rapid and effective decisions. Often, the decisions are made sequentially, and the consequences are likely to affect future choices. Examples of such operations are found in a variety of military and non-military systems.

Among these are anti-submarine warfare (ASW), command and control, shipboard tactical operations, coordination of electronic warfare (EW), control of remotely piloted vehicles, environmental surveillance, crime prevention, and air and highway traffic control. All of these applications have in common changeable decision environments, frequent decision responses, copious but fallible information, and a minimum of time available for off-line aids.

The human decision maker (DM) generally performs sub-optimally under such conditions. Cognitive limitations on memory, attention and processing, and biases and inconsistencies in aggregating information typify his behavior. Accordingly, computer aiding techniques have been developed to improve performance in many of these areas. Computer aiding can unburden the operator of routine computational tasks, assist in structuring the decision, perform probability assessments, and call attention to critical events. In general, these techniques have been applied to static, well-defined decision situations, using off-line modes of interaction.

More appropriate for dynamic tasks, however, where complex future consequences must be considered, is the use of the computer to observe and respond to the human during the decision making process itself. This type of interactive participation requires computer adaptation to changing task requirements and operator needs. Also, the often subjective, and incompletely quantified nature of real world decisions necessitates the incorporation of some form of a adaptable model of the human decision maker, in order to determine his preferences and goals.

1.3 System Concept

A decision support system termed ADDAM (Adaptive Dynamic Decision Aiding Methodology) was developed in accord with these guidelines. ADDAM consists of an adaptive decision model which continuously observes both the decision environment and the DM's behavior, learns his decision policy, and

makes choice suggestions based on the apparent value of the alternatives to the decision maker.

The adaptiveness of the ADDAM system is realized through the use of a trainable multi-category pattern classifier. As the DM performs the decision task, this on-line estimator observes the operator's choices among the various decision options. The estimator, using event probabilities as inputs, attempts to classify these probability patterns by adjusting utility weights according to an adaptive error correcting algorithm. In this manner, the utility estimator tracks the operator's decision making and learns his utilities. Such an approach has a number of advantages compared to off-line utility estimation. Dynamic estimation observes and models actual behavior rather than responses to hypothetical decisions. It does not interrupt or intrude on the process of decision making. And it responds to ongoing changes in task characteristics and operator needs.

In the current implementation, ADDAM is used to augment operator performance in two types of related decisions. First, the operator decides on a means of information acquisition. Second, on the basis of information received, he selects an appropriate action alternative. The action decision introduces dynamic considerations, because it impacts future information acquisitions. The ADDAM system aids in the information acquisition phase by inferring the operator's utility structure, combining these utilities with estimates of information availability, and recommending the information source with the highest expected utility. Complementary aiding in the action selection phase is given by a probability updating program. Revision of probabilities following information acquisition is computed using a Bayesian approach.

1.4 Experimental Studies

Several experimental studies were conducted to evaluate the decision support system in realistic but controlled circumstances. Prior to the study reported herein, a fishing fleet simulation task was used as the experimental vehicle. This simulation involved tracking and reporting the location of several components of a fishing fleet as they moved over an ocean expanse. The experiments focused on basic system validation and on how aiding affected the "internal" quality of decision making. The experimental evidence gathered from these studies indicated that (1) the adaptive model accurately predicted the operator's decisions, (2) aiding significantly improved decision consistency, (3) aiding significantly improved decision quality, (4) aiding reduced intersubject variability, and (5) aiding increased decision throughput.

The present study showed that "external" measures of decision outcome, such as accuracy, errors, etc., are also improved by adaptive aiding. The experimental task was a new ASW version of the fishing fleet simulation. Operators tracked the movements of a submarine and an interfering object, using the same types of sensors (sonobuoys, helicopters, MAD's, etc.) available in Naval ASW exercises. The sensors varied in reliability, specificity, and cost. Unaided operators worked alone, using computer-generated intelligence reports. Aided operators received additional computer assistance in the form of (1) sensor output evaluation, and (2) sensor placement recommendations. The latter were based on their own utilities for sensor information, derived in previous training sessions by means of a trainable pattern recognition program.

Twelve operator subjects participated in the study. They came from the local Air National Guard center, were equally divided between CO's and NCO's, and represented the type of military personnel who might interact with computer-aided command systems. Half were assigned to the aided group

and half to the unaided (control) group. Following two practice sessions, performance was recorded in a 1-1/2 hour test session. The basic performance score was defined as:

$$Score = Gain - Cost$$

where

$$Gain = Points - Penalties$$

Points were credited for correct submarine location reports, and *Penalties* were deducted for incorrect ones. *Cost* was the cost of sensor resources allocated. Operators attempted to maximize their score; score feedback showed them how they were doing throughout the test session.

Results showed that the aided group performed significantly better than the control group ($P<.05$), improving their mean score by almost a factor of two (88%). Improvement was partially attributed to a small but significant increase in the number of decision trials completed during the session. But most of it appeared due to the better overall quality of the aided decisions. That is, the aided operators incurred slightly higher costs, but received a much greater return in points, and a substantially lower number of penalties. Decision consistency, as measured by mean deviation from maximum expected utility, was significantly enhanced for the aided group, as in previous studies. And also in replication of previous studies, the improved performance of the aided group was accompanied by decreased intragroup variability.

1.5 Conclusions

The combination of probability aggregation, adaptive utility modeling, and normative decision recommendation appears to be well suited to the complexities of dynamic decision aiding. As implemented and evaluated in the ADDAM system, these techniques respond to differences in decision style

and to changes in task circumstances, they supply a variety of aiding information, and they provide a framework for analyzing and communicating the rationale for decisions to the decision maker. All in all, the adaptive system becomes an interactive partner rather than an inflexible set of programmed responses or an autonomous surrogate.

Much of the significance of the present work lies in its incorporation of both descriptive and normative aspects of decision analysis into a single system. The descriptive modeling is internally validating, through prediction of behavior, while the normative recommendations represent a processing of new inputs according to the same previously observed behavior. Thus, the operator's decision policy can not only be captured and analyzed, but can be used to systematically unburden him of some of the cognitive effort of processing new problems. Operator consistency tends to improve, since the aiding system incorporates a more representative sample of behavior than the person normally considers, and actions are structured more carefully. Finally, the technique is computationally parsimonious. Only those structural aspects necessary for capturing the operator's behavior are included, much as pattern recognition techniques normally require only a small portion of the structure required by complete dynamic models.

Analysis indicates that the domain of greatest promise for an adaptive decision support system lies in areas of high time and load stress, and where there is a reliance on changing, subjectively determined criteria of performance. Also, a basic structuring of the decision task, and some degree of recurrent behavior, is necessary for the continuing estimation of model parameters. In circumstances that do not satisfy these conditions, adaptive systems can work in conjunction with a variety of other analysis techniques, functioning as part of a decision support system similar to that envisioned by Leavit, Alden, Erickson and Heaton (1974). For instance, long range planning functions may be best accomplished using tree elicitation and structuring (Brown, Hoblitzell,

Peterson and Ulvila, 1974; Nickerson and Feehrer, 1975), while completely formulated and specified problems benefit strongly from the rapidity of decisions rules or preprogrammed responses (Brown, et al, 1974).

2. THE PSYCHOLOGICAL BASIS OF ADAPTIVE AIDING

2.1 General

This chapter provides a general background relevant to man's capability as an information processor and decision maker. In particular, the results and implications of empirical studies of his information gathering behavior are highlighted. This analysis leads to the rationale for developing procedures to aid man in the selection and utilization of information. Several decision aiding techniques are reviewed which have been based on procedures involving probability aggregation, utility estimation and adaptive models. In addition, some studies are cited which have assessed user acceptance of experimental computer aids. The chapter ends with some conclusions concerning the nature and characteristics of the particular decision aiding approach developed in the present program.

2.2 Man as a Suboptimal Decision Maker

The deficiencies and limitations of human performance in complex information processing and decision making tasks are well documented as a result of many years of behavioral research. Comprehensive reviews of the experimental literature are available, for example, in Lee (1971), in Slovic and Lichtenstein (1971), in Rapoport and Walsten (1972), in Nickerson and Feehrer (1975), and in Slovic, Fischhoff, and Lichtenstein (in press). Most of the experiments reviewed have had the goal of (1) describing human decision behavior and (2) gaining a better understanding of the cognitive processes humans employ to solve decision-related tasks. The specific tasks studied have generally been well-structured, to permit optimal or prescriptive models (e.g., Bayesian, regression, dynamic programming) to be applied to the same task parameters or input data which are presented to the experimental subjects. By comparing the output provided by the subject, who is presumably employing some sort of processing model, with the output

of the prescriptive model, the investigator can determine how reliably, and to what degree, human judgments match or depart from optimal or normative judgments.

The most general finding from behavioral experiments that have used the man vs. model paradigm is that the intuitive responses of the unaided human mind usually depart from those generated by an optimal model, but they do so in a systematic and stereotypical fashion. Thus, borrowing engineering terms to describe human behavior, man has been referred to as a suboptimal or inefficient information processor and decision maker. For example, when purchasing information, he may acquire a significantly greater or lesser amount of information -- depending on the particular situation or circumstances -- that called for by the normative model. Or, when making diagnostic inferences from observed data, man extracts either a significantly greater or lesser amount of diagnosticity -- again, depending on the specific situation -- than warranted by the optimal model. In the present program, we were particularly interested in information gathering, and in the opportunities for decision aiding of this process.

2.3 Information Gathering Behavior

2.3.1 Description. In their recent comprehensive review of the information processing and decision making leterature, Nickerson and Feehrer (1975) identify "information gathering" as one of the principal tasks to be performed in a decision-oriented system. They describe the process as follows:

> *"From the point of view of the decision maker, most decision situations are characterized by some degree of uncertainty. This uncertainty may involve the current "state of the world," the decision alternatives that are available, the possible consequences of selecting any*

given one of them, and even the decision maker's preferences with respect to the possible decision outcomes. One of the major problems facing the decision maker, therefore, is that of acquiring information in order to reduce his uncertainty concerning such factors, thereby increasing his chances of making a decision that will have a desirable outcome.

What makes the problem interesting, and nontrivial, is the fact that information acquisition can be costly, both in terms of time and money. Therefore, the decision maker must determine whether the value of the information that could be obtained through any given data-collection effort is likely to be greater than the cost of obtaining it. And therein lies a decision problem in its own right."

2.3.2 Experimental Studies. Much research has been done on information acquisition in decision tasks, but most studies have concentrated on information purchasing behavior. The latter have, for the most part, failed to capture the complexity of the problem that often faces the information seeker in the real world. In the typical information purchasing experiment, information from a single source is presented to the subject, with his task essentially to decide whether it's worth what it will cost to acquire it. However, in many practical situations, the decision maker or commander must seek and locate the information he needs or wants, and he must often select from among alternatively available information sources of differential quality in terms of information diagnosticity and cost. A few relevant studies of information selection behavior are reviewed below.

Kanarick, Huntington and Peterson (1969) studied performance in a simulated scenario in which the subject had to reach a binary decision about whether an enemy submarine was either present or absent in a given vicinity.

On each trial the subject could purchase data, from one of three different information sources. The sources varied in both reliability and cost, where the higher the diagnosticity of source, the greater the cost for consulting it. The penalties for incorrect decisions were also manipulated experimentally. Subjects' behavior was sensitive to the variations in the independent variables; however, performance was deficient when compared with an optimal Bayesian model. For example, they consulted the most reliable (and most costly) sources less frequently and the less reliable (less costly) sources more frequently than they should have. Also they generally purchased less information than required by the optimal model. This last result might be accounted for by the common finding of recent research that shows subjects tend to over-estimate the diagnostic impact of less than perfectly reliable data (e.g., Johnson, Cavanaugh, Spooner, and Samet, 1973).

Although Kanarick, Huntington and Pearson (1969) allowed subjects to choose among multiple information sources, the sources were presented in parallel. However, in real life situations, information from various sources is frequently sought, generated, and made available in a sequential, rather than in a parallel, mode. In a dynamic situation, furthermore, the uncertainty of the environment may force the information seeker to perform under fluctuations and restrictions in the amount of available information and/or the level of resources needed to acquire the information. For example, the military commander can take advantage of all the patrol units that he can spare but still require more information about the enemy.

To study information seeking behavior under these kinds of conditions, an experiment was conducted by Levine, Samet, and Brahlek (1975). These investigators required subjects to determine which of four populations was being sampled in a multinomial Bayesian task. Each sequentially drawn datum was described on one of three dimensions which represented different levels of information source diagnosticity (high, medium, and low). On each

trial, the subject purchased knowledge of the identity of the information source which was available on that trial, and he had the option to either purchase the associated datum at a fixed additional cost or pass it up at no additional cost. Using this paradigm, the amount of information potentially available and the percentage of it which could be purchased by the resources provided were varied factorially, and the effects on information selection and purchasing behavior was assessed.

The principal relevant findings were that: (a) relative to the low diagnostic source, subjects purchased information about 5 times as often from the medium and high diagnostic sources; (b) when more information was potentially available, subjects were more efficient -- relative to an optimal Bayesian model in selecting from among information sources; (c) significantly more information was sought as both amount of available information and purchasing resources increased; and (d) across all experimental conditions, subjects generally purchased more information than was recommended by the normative model. With regard to the last finding, the subject was actually paying for information which had a negative value, i.e., information whose acquisition led to a decrease in expected payoff.

In another study involving the selection of information, Rapoport, Lissitz, and McAllister (1972) investigated the search behavior of subjects required to find a single object hidden in one of four distinguishable locations. For each location, they were given: (a) the *a priori* probability that the object would be detected there, (b) the probability that the object would not be found by a search (i.e., a random sampling); and (c) the per-trial cost of search. In agreement with the previous studies, the results indicated that subjects do not consult information sources in an optimal fashion. Of particular interest was a finding suggesting individual differences in search strategies; that is, some of the subjects deviated from the optimal policy in the direction of maximizing detection probability, whereas others deviated in the direction of minimizing per trial search costs.

The results of experiments on information purchasing behavior relate to man's capability as an information selector. The typical experimental paradigm allows the subject on each trial the option of either purchasing more data relevant to the decision that he is required to make, or to stop data collection and make a decision. Stopping data collection is also a decision and defines the selection of a predecisional information set. The various studies have shown that subjects are highly sensitive to informational and situational parameters, e.g., environmental variance (Schroeder and Benbasat, 1975), *a priori* probabilities for decision alternatives (Green, Halpert, and Minas, 1964), data diagnosticity (Snapper and Peterson, 1971), source reliability (Levine and Samet, 1973), and costs and payoffs (Pitz and Reinhold, 1968; O'Connor, Peterson, and Palmer, 1972), but their performance departs systematically from optimal performance. In general, it appears that too little information is purchased when much is required by a Bayesian model and too much information is purchased when little is required. For example, subjects have been found to require from two to nine data observations to revise their opinions as much as Bayes' theorem would prescribe for one observation (Peterson, Schneider, and Miller, 1965; Phillips and Edwards, 1966).

2.3.3 Rationale for an Information Gathering Aid. It is evident that when information quantity is held constant, an improvement in information quality leads to an improvement in decision performance (e.g., Levine and Samet, 1973; Snapper and Peterson, 1971). With regard to military information processing systems, the issue has been stated as follows: "The key to competent decision making is the availability of current and accurate information. It is not the quantity of information which is important. Rather, it is the process of selecting the pertinent information, assessing its significance, and displaying it in a readily understood format which facilitates the decision making process" (Albright, 1975; Levit, Heaton, and Alden, 1975). Since one way to achieve an increase in information

quality is to be more selective in collecting information, we can ask how well man does as an "information selector" or discriminator among alternative information sources. The basic conclusion reached by each of the experiments reviewed above is that although subjects are sensitive to the differences in information source quality, they perform suboptimally in selecting among information sources.

If man is sensitive to key informational and situational parameters, why does he consistently show systematic, stereotypical biases when choosing among available information for decision making? Apparently, because of his limited memory, attention, reasoning, and computational capabilities. He is unable to integrate/aggregate/combine various dimensions of information -- each with its associated graded level -- to arrive at a compositive, subjective value for the information which is consistent and valid. Thus, for example, he is unable to appropriately trade off the reliability of information against its cost (Kanarick, Huntington and Peterson, 1969). Therefore, it appears that a valuable type of aid would be one which helps him to assess and to apply consistently his own utilities for the information provided by alternative sources. As shown later, this is the approach adopted in the ADDAM system.

2.4 Previous Decision Aiding Systems

2.4.1 Definition. Nickerson and Feehrer (1975) have clarified a useful distinction by suggesting that "performance deficiencies" are presumably correctable behaviors, whereas "performance limitations" are not. To illustrate, they give "the tendency of humans to be overly conservative in their application of probabilistic information to the evaluation of hypotheses" as an example of a "deficiency", and "the inability of most people to weigh more than some small number of factors, without some procedural help, in arriving at a preference among choice alternatives" as an example of a

"limitation". Further, with regard to the implications for training decision makers, they state: "deficiencies may be 'trained out'; basic limitations must be 'trained around'." Whatever the case, it is reasonable to hypothesize that man can benefit from high level aids designed to help him recognize his inadequecies, overcome his deficiencies, and circumvent his limitations -- in other words, to become a "smarter" or "better" decision maker.

A generally acceptable definition of a decision aid is hard to come by. After a comprehensive search through the literature on decision aids, Levit, Alden, Erickson and Heaton (1974) concluded that "... the working definition of decision aiding is dependent on the assumptions of the decision making framework from which it is derived". However, most decision aids are conceived with the same basic intention in mind: namely, to enable a decision maker working with the aid to perform measurably superior to what he would do working without the aid.

2.4.2 Probabilistic Information Processing. The purpose of decision aiding is to allocate the performance of decision functions between man and machine in a way which optimizes the use of their respective strengths (and weaknesses). Much of the work on decision aiding has centered around Probabilistic Information Processing (PIP) (Edwards, 1962; 1964) and similar systems for Bayesian information processing (e.g., Kaplan and Newman, 1966; Kelly and Peterson, 1971; Howell, 1967; Johnson and Halpin, 1972). Central to this approach is evidence that human decision makers are well suited to estimating the conditional probabilities that specific predictive (diagnostic) information will be observed when certain environmental states occur, and that they have difficulty in aggregating the probabilities into opinions (Edwards, 1964).

Bayesian information processing involves the use of Bayes' theorem to help the decision maker optimally revise his opinion in light of new information. The technique allocates data evaluation to the man and data aggregation to the computer. The concept requires that the person estimate the likelihood ratio of each data point, and transmit these to the computer, which utilizes Bayes' theorem to aggregate the likelihood ratios and make inferences. By following this procedure, it is possible to compensate for the man's inability to retain and combine separate data points into an overall conclusion.

Several PIP-type systems have been demonstrated in addition to Edwards', including one developed recently by Johnson and Halpin (1972), at the Army Research Institute. This system, a multi-stage computer-aided Bayesian inference system, is now being tested in tactical intelligence environments using intelligence officers as subjects.

Another PIP-type system for assisting man in making diagnostic decisions about reconnaissance data, developed and implemented at Ohio State University (Howell, 1967), has had considerable success. Howell says that improvements in diagnostic decisions of around 10-15% can be expected with automated aggregation. Improvements become particularly noticeable under conditions of time or load stress or low input fidelity. Kelly and Peterson (1971) have reported similar findings with their intelligence analysis system. Intelligence analysts were trained in the technique and subsequently were tested in analysis of realistic intelligence problems. Results have indicated that the PIP procedure increased the "efficiency with which probabilities or odds are revised in light of new information".

Thus, Bayesian decomposition techniques seem to represent a successful attempt to employ a man/computer team for making statistical inferences. One reason for their capability to improve the inferential process is that they

appear to compensate for inherent operator biases. After a comprehensive up-to-date review of PIP studies, Beach (1975) has concluded:

> *"It is fairly apparent from the literature ... that the use of subjective probabilities and Bayes' theorem in real world decision making, whether it involves a military decision, ... a medical diagnosis, or ... is potentially profitable. But while there has been a great deal of laboratory experimentation using Bayesian techniques, it is clear that more research needs to be done in more realistic settings."*

2.4.3 Utility-Based Aiding. Bayesian information processing is an example of how successful allocation of function can be made between a man and a computer in an inferential task of transforming a data set to a hypothesis set. However, the procedure does not cover the overall decision process, since it involves only the estimation of probabilities. Assessments of risks and gains must be performed within the framework of statistical decision theory. A logical expansion of the PIP concept thus involves extending the idea of compensation for operator biases in probability assessment to cover operator utility assessment.

A system which makes use of operator utilities was developed by Miller, Kaplan and Edwards (1967,1968). Using this system, the Judged Utility Decision Generator (JUDGE), they successfully demonstrated the advantage of automating aircraft dispatch in tactical air command systems. In this situation, the commander was required to consider the relative value of the targets, their probability of destruction and the available aircraft. In forming a decision policy, the key human inputs were the utility of destruction of various targets. JUDGE used the value judgments of trained personnel as input to the computer. Other inputs to the computer included demand forecasts, number of aircraft and turnaround-time

distribution. The computer then selected a course of action which maximized the expected utility. Experiments with experienced tactical air controllers revealed that JUDGE did a more effective job than did men. Of particular significance in this work was the method of eliciting utilities and checking these utilities for logical consistency, coherence, validity and reliability (Newman, 1975).

The JUDGE system represents a class of automated decision systems. This type of system incorporates the operator's utility and subjective probability judgments into a maximum subjective expected utility (SEU) model. The rationale is that man observes the real world, estimates its states and their associated probabilities and utilities and provides them as inputs to a computer algorithm which then generates the decision output (Schum, 1970).

JUDGE extended decision aiding into the area of utility assessment. Another area for extension is the automated estimation of probabilities. PIP-type decision aiding systems place heavy reliance on human "transducers" to generate their input parameters (conditional probabilities). As the number of parameters increases, this becomes more and more difficult. Likewise, it becomes difficult for humans to estimate conditional probabilities in non-stationary environments. Howell (1967) felt that the current state-of-the-art did not permit the automation of the conditional probability estimation process. Freedy, Hull, Lucaccini and Lyman (1971), however, have demonstrated the feasibility of automating the estimation of conditional probabilities in certain stationary and non-stationary environments.

2.4.4 Adaptive Decision Models. Adaptive models for decision aiding attempt to learn the decision process of the human operator by (a) successive observation of his actions, and (b) establishing and interim relationship between the input data set and the output decisions (the

model). Learning in this context refers to a training process for adjusting model parameters according to a criteria function. The object is to improve model performance as a function of experience or to match the model characteristics to that of the operator. There have been a number of approaches to trainable systems involving operator modeling, operator aiding and replacement of the operator by his model.

Research of major interest in this area is the "bootstrapping" procedure based on an autocorrelation model of the operator (Bowman, 1963; Dawes, 1971; Goldberg, 1970). The parameters of the decision model are adjusted by a regression process to obtain a "least squares" fit to the actual decision policies of the operator. Data for the parameter adjustment are obtained from actual data collected from the operator's previous judgments. The model then can be used to replace the decision maker in an identical decision process. The bootstrapping technique is inherently limited in its ability to handle decisions with large dimensions of values (Fischer, 1972) and by its *a posteriori* application. However, the result of experimental work with real world problems reinforces the assertion that the operator can be effectively replaced by a model derived essentially from his own judgments.

Additional research of interest was done by Rouse (1972). Rouse used a linear quadratic learning function for modeling operator behavior in a dynamic prediction task. This model was also based on a regression and autocorrelation approach to learning. Its primary purpose was to determine the effects of constraints such as memory limitations on model performance. The use of a learning system to establish an adaptive model of the human operator through real time parameter tracking was also reported by Gilstad and Fu (1970). Although the approach was applied in the context of two-dimensional compensatory tracking, it illustrates how a learning process can be used to identify decision and control behavior. Linear and piecewise linear discriminant functions were used to classify

system gains, errors, and error rate. The decision boundaries for classification were determined through a process of on-line learning -- observing operator performance and parameter adjustment. The specific model used was applicable only to very limited tasks, and merely illustrated the feasibility of the technique.

An interactive adaptive system, termed the Autonomous Control Subsystem (ACS) was used to share decision making functions with a human operator. The ACS was originally used to control a remote manipulator (Freedy, et al, 1971) and was subsequently implemented in a simulated, continuous, decision and control task (Freedy, Steeb, and Weltman, 1972). Studies using ACS showed that human interaction with an "intelligent" machine is very different from interaction with a deterministic machine. Deterministic decisions are generally accepted as correct within their context. The intelligent computer, on the other hand, introduces the idea of machine fallibility. Presented with apparently equivalent data, automata can remain indecisive, jump to conclusions, commit errors, learn slowly or rapidly at different times, and perhaps even develop behavioral quirks. Accordingly, optimal allocation of system function depends on continuous and accurate assessment of the "mind of the machine" by the operator or by some external monitor.

It has been found that adaptive models can perform better than the operators whom they model. That is, the correlation between the output of the model with an appropriate criterion (e.g., correct decisions) is often higher than the correlation between the operator's output and the criterion. This important finding has been demonstrated in various real world contexts: managerial decision making (Bowman, 1963), prediction of academic success (Dawes, 1971; Wiggins and Kohen, 1971), and clinical diagnosis (Goldberg, 1970).

Why do models or decision rules based on an operator's average behavior pattern outperform the actual decision behavior of the operator? Several investigators have suggested that the superior performance of the models is due to their ability to eliminate or reduce "noise" effects in subjective weighting of evidence and in erratic operator responses. For example, Bowman (1963) described the filtering process as follows:

> *"Man seems to respond to selective cues in his environment ... particular things seem to catch his attention at times ... (These random and particularistic components can be eliminated) through the use of decision rules incorporating coefficients derived from (the operator's) own recurrent behavior."*

Dawes (1971) put it somewhat differently by stating that:

> *"A mathematical model, by its very nature, is an abstraction of the process it models; hence, if the decision maker's behavior involves following valid principles but following them poorly, these valid principles will be abstracted by the model -- as long as the deviations from these principles are not systematically related to the variables the decision maker is considering."*

However, after a critical review and analysis of linear models in decision making, Dawes and Corrigan (1974) concluded that the success of these models is tied up with their inherent robustness and with the fact that they have been tested largely in situations having particular characteristics most appropriate to their application.

2.5 Acceptance of Computer Aiding

There is growing concern with the need to "humanize" computer-based information systems in order to increase their acceptance by users (Sterling, 1975). Several studies have looked into the use of computer decision aids and the problem of man/computer interaction. Hanes and Gebhard (1966), in a realistic simulation of anti-air warfare, found that most naval commanders freely accept computer advice in tactical command action. Acceptance was determined by the decision logic, the nature of the tactical situation, the rate of presentation of recommendations, and the nature of the man/computer communications and display systems.

It is generally agreed that human decision makers must understand the capabilities and limitations of their decision aids. Unless they have this understanding, the computer aids are misused and overall system performance deteriorates, particularly as the importance of the decisions increase (Schaffer, 1965; Samet, 1969; Myers, Gibb, and McConville, 1963). One factor which contributes to the DM's attitude toward computer advice is his opinion of its performance accuracy. Halpin, Thornberry, and Streufert (1973) conducted an experiment using a "staff" computer which performed both accurately and inaccurately. They found, as might be expected, increased utilization for increased accuracy. They also found that exposure to an initially inaccurate computer aid decreased later utilization of an accurate aid (Halpin, Johnson, and Thornberry, 1973).

2.6 Present Adaptive Aiding Concept

The preceding sections have presented an overview of important research relating to computer aiding, and have pointed out a number of areas of potential practicality. It should be remembered that the present research program is directed toward the practical, rather than the theoretical

aspects of decision aiding -- i.e., toward developing usable on-line adaptive decision aids and studying the human factors involved in their use, rather than developing more accurate psychological models of human behavior. If more accurate models are developed as a result, however, they will be a welcome by-product.

The adaptive aiding system described in the following chapters combines the following important aiding approaches:

(1) *Aggregation of Decision Information*. The value of Bayesian aggregation techniques for dynamic decision aiding has been demonstrated in PIP-type systems. They are used here for estimating the probabilities of real world state transitions. In addition, computational assistance in assessing risks and gains, by means of an expected utility model, is provided. The approach differs from the automated decision systems (e.g., JUDGE) in several respects. Its purpose is to assess multiple alternatives and suggest high value alternatives for human evaluation, rather than to automate the decision process. Its output is not intended to be the optimal solution to a static decision task.

(2) *On-Line Adaptive Estimation of Decision Parameters*. Existing aggregation systems place heavy reliance on human estimators of decision parameters (conditional probabilities and utilities). The Perceptronics approach uses on-line adaptive estimation of decision parameters for several reasons: (a) dynamic decision tasks inherently require on-line aiding techniques

because they involve sequences of decision based on the results of earlier decisions; (b) adaptive estimation of decision parameters relieves the DM of this task; (c) it helps eliminate DM biases, noise perceptions, and memory limitations; and (d) it can respond to non-stationary environments. Adaptive parameter estimation allows the DM to devote his attentions to the evaluative aspects of decision making.

(3) *Normative Models of Decision Behavior*. A mathematically simple normative model for human decision making behavior (an expected utility model) is used. Its primary purpose is to help form a symbiotic relationship between the human decision maker and the aiding system. The model adjusts its parameters (utilities) in order to track the DM's decisions. It helps stabilize his behavior when the information sources conflict or have low reliability and calls his attention to high value alternatives which he might otherwise overlook.

2.7 Significance of ADDAM as a Decision Aid

The domain of application for the decision aid developed here is a rapidly changing environment, where decisions of the same general type must be made repeatedly. The modeling technique is based on the prediction of decision behavior according to an expected utility (EU) model. If the man and the computer can be considered as representing a single system, then the goal of the technique is to improve the decision output of the system.

By focusing on the underlying principle employed by ADDAM to model behavior, insight can be gained on its significance as a decision aid.

ADDAM adjusts model parameters only after an incorrect prediction, i.e., one that is not consistent with the previously observed decision behavior of the operator. Since incorrect prediction is the key factor which controls and determines the estimation of the model parameters, it makes sense to examine closely what appear to be the possible sources of incorrect prediction. These are model inappropriateness, parameter change, and random error. Each will be discussed separately, with particular emphasis on how it impacts upon the psychological basis of the ADDAM technique.

Model Inappropriateness. One possible source of an incorrect prediction is simply that the EU model on which ADDAM is based is not appropriate for describing the decision behavior of a human operator. That is, the operator may not combine probabilities and utilities in accordance with a normative model and/or may not necessarily choose a course of action to maximize the EU of the outcome. It is also possible that the form of the model used by the operator is generally appropriate but that he does not take into account all the constituents (i.e., outcome combinations) which are normatively required. Whatever the case, despite challenges to the feasibility of EU as a behavioral model (e.g., Barron, 1975), EU persists as a convincing and general explanation for decision making in structured situations (Rapoport and Wallsten, 1972).

Within ADDAM, utilities are inferred from behavior rather than directly elicited from the operator, and it is not necessary to assume that the operator consciously calculates EU. Furthermore, when ADDAM recommends decisions to the operator which maximize his EU, it has the capability to shape his behavior pattern in accordance with an EU model. On a psychological basis, this conditioning approach toward making decision behavior more rational appears more promising than a direct, instructional approach which has proved unsuccessful (Lichtenstein, Slovic, and Zink, 1969).

Parameter Change. If it can be assumed that the EU model is generally appropriate for a particular decision task situation, then the "goodness" of prediction depends upon how closely the model parameters, i.e., utility weights, approximate the true utilities currently manifested by the operator. Since the operator is performing in a dynamic, changing environment, it is reasonable to expect that his utilities for specific outcomes will vary with time and the situation. Thus, another possible source of an incorrect model prediction is a change in parameter values. That is, if a prediction based on a set of utilities is incorrect, it could be that the utilities have changed so that the set on which the prediction was made are not longer valid.

Since ADDAM is on-line and dynamic, by definition it can sense and keep track of systematic changes in utilities. To improve the accuracy of prediction, ADDAM repeatedly makes adjustments in the utility weights with a minimum of time lag. This adaptive ability of ADDAM represents one of its strongest advantages, especially over static utility estimation techniques, in that utility-based decision aiding can incorporate the most up-to-date estimates of operator utilities.

Random Error. The third source which can contribute to an incorrect model prediction is random error. By this we do not mean "noise" in the model prediction mechanism but rather random error in the responses of the operator. If it can be assumed that the EU model does adequately describe the operator's decision behavior and that his utilities are correctly estimated, then why should his decisions be discrepant with those predicted by the model? The answer lies in the fact that man is not a perfect information processor and decision maker; he is sometimes inconsistent or even erratic due to "noise" effects generated by the environment and his interaction with it.

One way to decrease the opportunity for an operator to make a decision which is inconsistent with his value structure would be to provide him, prior to his making a decision in a given situation, with the recommended decision which is consistent with his behavior as modeled to that point. This is precisely what ADDAM is designed to do and it is in this feature where its unique significance as a decision aid is most pronounced. By contributing to the reduction of random error in the human component of the man-machine decision system, ADDAM has the potential to improve the "cognitive reliability" (Halpin, Johnson, and Thornberry, 1973) of the operator and thereby augment significantly the quality of the system output.

In summary, the EU model on which ADDAM is based is robust, and has been successfully demonstrated as a useful procedure for modeling human decision behavior. Since the model is on-line with the operator, it can take into account justifiable fluctuations in the operator's utility function, and adjust itself accordingly. By prompting the operator with model-recommended decisions, the aiding system helps him to act as a more consistent and rational decision maker. As a result, it can improve the overall performance of the man-machine decision system.

3. ADAPTIVE AIDING METHODOLOGY

3.1 ADDAM System Overview

Adaptive or goal directed techniques are employed extensively in the ADDAM decision support system. The adaptive methods involve the on-line acquisition of operator decision strategies by computer observation of his behavior. This dynamic modeling is capable of in-task observation of operator decisions made in response to real world probability data. The decision maker's value structure is then computationally inferred through a pattern recognition algorithm, and used as an input to a decision recommendation program. The resulting behavioral model and aid have the advantages of (1) functioning operationally in actual tactical circumstances, (2) adapting to changing task requirements and operator capabilities, and (3) requiring minimal programming complexity.

The adaptive modeling described here is related closely to the "on-line model matching" methods practiced in adaptive manual control (Gilstad and Fu, 1970) and to the adaptive linear models used to augment or replace the expert decision maker (Bowman, 1963; Kunreuther, 1969; Dawes and Corrigan, 1974). These techniques use pattern recognition or learning algorithms to estimate behavioral parameters. The ensuing models are then used to train, replace, or evaluate the operator. The current work extends this field of work by placing the operator in a real-time interaction with his model. The system both descriptively models and proscriptively aids the operator.

The complete ADDAM Decision Support system covers a variety of decision processes. The system includes programs for probability aggregation and updating, utility estimation, and strategy recommendation. The modes of decision support include: (1) revision of probability estimates upon receipt

of new information, (2) predicting the consequences of selected alternatives, (3) suggesting choices which optimize the operator's values, (4) providing a basis for comparing behavior to organizational standards, and (5) communicating the rationale for decisions to others.

Because the decision model is adaptive, model based decision aiding establishes a complex symbiotic relationship between the operator and ADDAM. The system adapts to the human operator's pattern of behavior and, in turn, provides decision aiding which may cause the human to modify his behavior. In a sense, the decision maker is provided a tool with which to refine his behavior. Rather than confronting each decision anew, and depending on often fallible processes of recall, recognition, problem structuring, and evaluation, the operator uses logically derived recommendations to guide and condition his responses.

Three major elements are required to apply the on-line model matching approach to dynamic decision situations: (1) a theoretical structure for modeling the decision process, (2) model learning and parameter estimation techniques, and (3) performance evaluation techniques which can appraise the quality of both the internal decision process and the external performance. The development of each of these elements is detailed in the following sections.

3.2 ADDAM Theoretical Structure

The ADDAM Decision Support System is composed of a combination of three complementary elements -- a set of probability aggregation programs, a dynamic model for tracking operator values for outcomes, and a strategy recommendation algorithm. Each of these aiding subsystems has a major role in augmenting the human functions of problem formulation, analysis, resolution and evaluation.

3.2.1 Probability Aggregation. The first mode of aiding, probability aggregation, is possibly the most established and procedurally defined area of support. This type of aiding is typically based on Bayes' rule, a mathematically appropriate way to revise probability estimates with new information (Edwards, 1962; Johnson and Halpin, 1972; Beach, 1975). The technique combines prior probabilities and conditional probability estimates to arrive at posterior probabilities. Bayesian aggregation is used by the ADDAM system to update environmental status and sensor data when new information is received.

3.2.2 Utility Estimation. More difficult are the considerations of perceived gains associated with the decision outcomes. Occasionally, objective values in terms of dollars, ship-equivalents or other external criteria can be used as criteria for choice. The situation must be exhaustively quantified to justify this type of calculation. For instance, a strategy for action selection based on such objective criteria such as speed, accuracy or expected value may be relatively easy to derive when system objectives, behavior, and environmental conditions are completely specified. Given the immediate probabilities of obtaining the possible outcomes and given the costs of the consequences, the decision choice with the highest expected value can be selected. Objective performance criteria for the immediate task in most man/machine systems, however, are not well defined, or are only indirectly related to long term system goals. This indeterminacy is particularly evident in systems operating in dynamic environments, where the results of earlier decisions affect later decisions. Such systems may rely heavily on the operator's subjective evaluation of the situation at hand, and the decisions should be based on measurable subjective preferences (utilities) of the operator.

Numerous techniques are available for assessment of the operator's utilities, ranging from ad hoc procedures to completely axiomatic analysis.

The simplest techniques entail eliciting direct expressions of preference along qualitative or quantitative scales. Fishburn (1967) lists more than a dozen such direct methods. Other techniques of utility assessment include the decomposition of complex decisions into hypothetical lotteries, and the use of multivariate methods to analyze large numbers of binary preference expressions to determine underlying factors (Kneppreth, Gustafson, Johnson, and Leifer, 1974).

A major practical limitation to the application of decision theory is the complexity of utility assessment techniques. Most applications require a two-step process. The first step is to assess the decision maker's (DM) utilities, and the second is to apply them to the decision problem. Because it is not feasible to reassess utilities frequently in repetitive tasks, it is assumed that they remain static during this application. Such an assumption might be valid for a "one-shot" decision. However, there is no reason to assume that the DM's utilities remain static during the performance of multi-stage decision tasks. Nor is it reasonable to assume that they remain the same when the context changes from that of a laboratory context to the real world task.

The technique developed in ADDAM for dynamic utility estimation circumvents many of these problems. Dynamic estimation uses the principle of a trainable multi-category pattern classifier to "learn" the operator's utilities for the outcomes of information acquisition decisions (Freedy, Weisbrod, and Weltman; 1973). Such an application of pattern classification techniques was first suggested by Slagle (1971), who pointed out that the utility function was an evaluation function which could be learned from a person's preferences. The adaptive technique, described more fully in Section 3.3, assumes an expected utility maximization paradigm for modeling decision behavior, and uses a pattern recognition algorithm to successively adjust the model to fit observed decision behavior. The underlying expected

utility (EU) model assumes that the operator chooses that action whose expected (probability weighted) utility of outcome, is highest (Krantz, Luce, Suppes and Tversky, 1971). EU models, of course, are not a panacea for structuring decision models. Lichtenstein and Slovic (1971) argue that descriptive models must take cognitive factors into account, Luce and Suppes (1965) question the use of deterministically maximized choices rather than stochastic choices, and Wendt (1970) and Coombs and Pruitt (1960) contend that the EU model should be modified to account for preferences in variance of outcome. In general, though, the usefulness of EU model is conceded in situations where the number of choices is low and the decision maker can relate to all attributes in terms of probabilities (Goodman, Saltzman, Edwards, and Krantz, 1971). Also, the EU models have the advantage of modeling both descriptive and normative (optimal) behavior, unlike most of the heuristic-based models (Wendt, 1973).

A comparison of the positive attributes of the major utility assessment techniques -- direct elicitation, gambles, and dynamic observation -- is illustrated in Table 3-1. This table is adapted from Kneppreth, Gustafson, Leiger and Johnson (1974). The advantages of the dynamic observation technique are as follows: (1) Utilities are estimated non-verbally, without the need for a skilled analyst highly trained in utility estimation techniques. In fact, 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 to changes in values and the utilities are automatically validated by direct comparison with the DM's real-world behavior.

TABLE 3-1. 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
COMBINATOR	AFTER ANALYSIS	YES	YES
ADAPTIVE TO VALUE CHANGES	NO	NO	YES

3.2.3 Strategy Recommendation. The third element of the ADDAM system, the strategy recommendation program, follows naturally from the probability and utility estimators. With these parameters defined it is a simple matter to recommend individually optimal decisions. The choice with the greatest expected utility is determined and displayed to the operator. The recommendations given are thus based on the operator's own apparent values, and are organized into a normative framework. A certain generality is present in the normative processing, since the recommendations are not restricted to the identical circumstances of the observations used for training. Recurrent observations of the operator actions and circumstances are necessary for estimation of parameters, but these determinations generalize to other circumstances of the same structure.

3.2.4 Aiding Dynamics. The strategy recommendation algorithm closes a man-computer decision cycle or loop of considerable flexibility and dynamics. The extent of the aiding can be observed by examining the major decision processes of information acquisition and action selection. These processes are diagrammed in Figures 3-1 and 3-2, without aiding. In the information acquisition task, Figure 3-1, the operator receives feedback of the data requested and of the costs of data acquisition. To achieve long term success, he must ascertain what type of behavior led to maximum performance, a difficult task with probabilistically unreliable information sources. He must then use the data obtained to select timely actions (Figure 3-2) and to evaluate his performance using sporadic or noisy performance feedback. This cycle repeats itself as information is converted into action throughout tactical decision making, and because of the dynamic nature of these cycles, errors tend to compound.

Aiding in Information Acquisition. The ADDAM system aids the operator in the information acquisition process by introducing several additional loops. Figure 3-3 shows the augmented structure of information acquisition

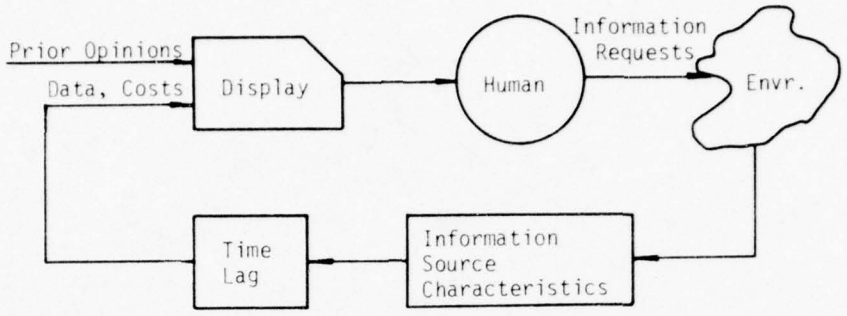


FIGURE 3-1. INFORMATION ACQUISITION

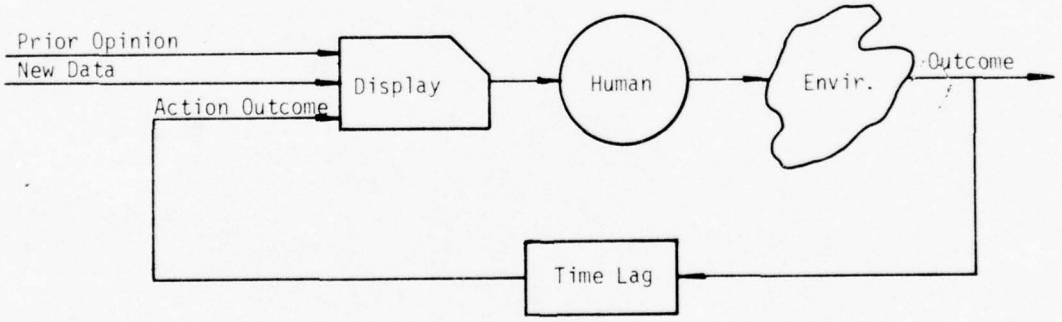


FIGURE 3-2. ACTION SELECTION

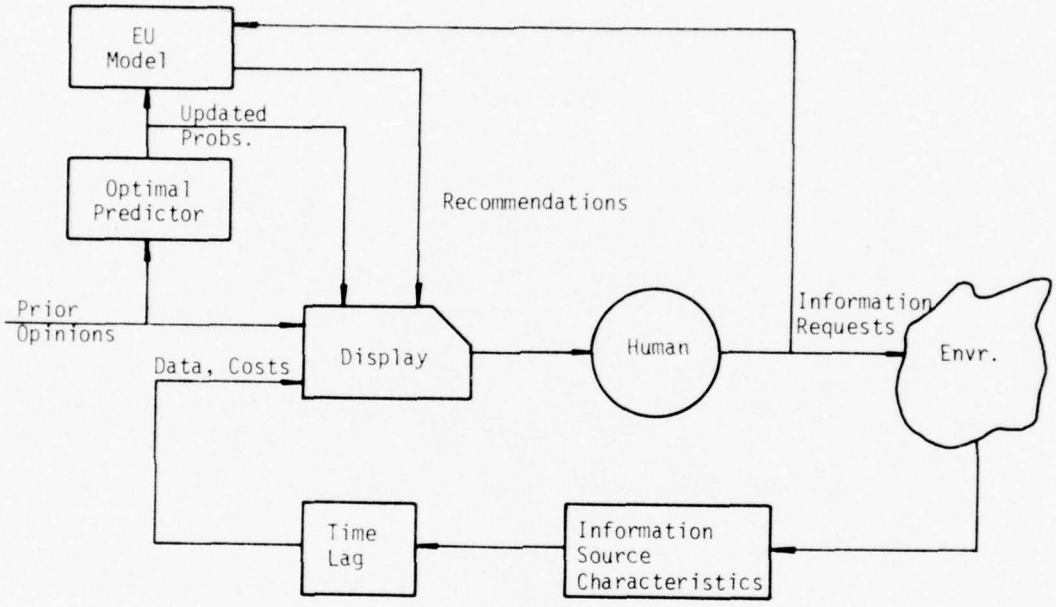


FIGURE 3-3. INFORMATION ACQUISITION WITH AIDING

process, demonstrating both feedforward and feedback loops. The feedforward part is a predictor display. It uses a model of the state transition probabilities to arrive at revised probability estimates of the state of the environment. These probabilities are both displayed to the person and used as inputs to the utility estimation program.

The EU model, comprising the lower feedback loop in Figure 3-3, is the heart of the ADDAM system. The model uses the predicted probabilities of states as inputs, and attempts to match operator behavior by adjusting weights decision-by-decision. This model is then used to recommend decision choices in subsequent decisions. These recommendations are not merely the result of calling up previously observed responses. The values are condensed from a variety of observed behaviors and in the recommendations are applied to decision circumstances that may be new to the decision maker.

Aiding in Action Selection. The second major decision process considered in ADDAM, action selection, is similarly amenable to dynamic aiding. Here, however, the aiding is entirely of the feedforward type, as shown in Figure 3-4. Probability aggregation and normative strategy recommendations are again made, but in different ways than in the information acquisition decision. Rather than updating probabilities of outcomes according to transition relations, the prior opinions are Bayesian updated according to the information received. This is a well established and completely deterministic calculation (Rapaport and Wallsten, 1972). It is also assumed that the outcomes of the actions are associated with set objective costs. The revised opinion can then be weighted by the objectivity defined costs for the outcomes and a maximum expected value choice recommended. This is a preprogrammed calculation, and does not require estimation of utilities or use of adaptive processes. If the outcomes could not be objectively quantified, a dynamic utility assessment and EU based recommendation algorithm of the type described earlier could be used.

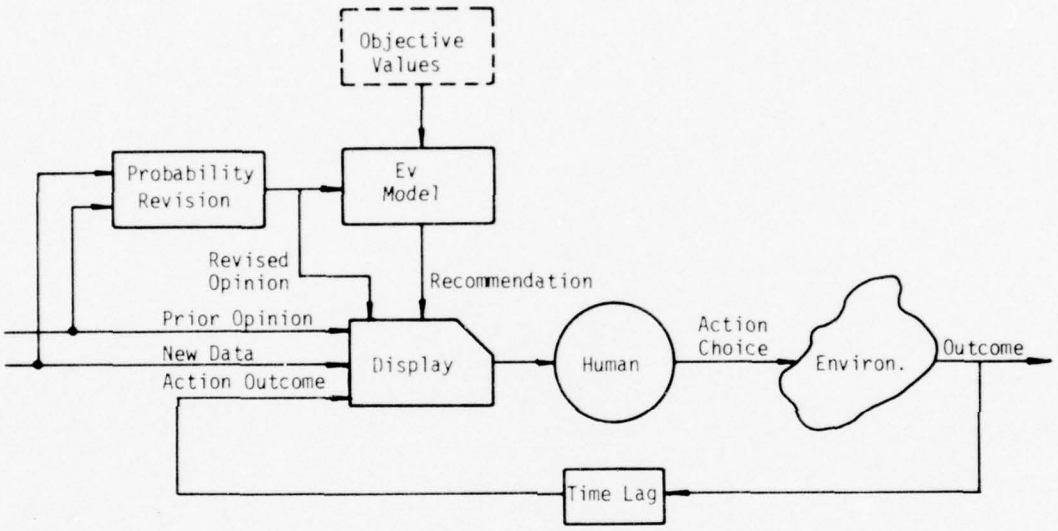


FIGURE 3-4. ACTION SELECTION WITH AIDING

3.2.5 Theoretical Evaluation. Use of the computer to both model the operator's values and recommend actions based on a normative decision rule using these values combines the best qualities of man and machine. Rather than relying on the limited abilities of man, or on the compromises inherent in a complete model replacement of man (as in bootstrapping), the system delegates the computational tasks to the machine and the evaluation and supervision functions to the man. In essence, the system applies the recommendation of Kunreuther (1969): have the machine utilize quantifiable data based on past behavior to give advice and have the human modify the machine recommendations according to external cues and system structural changes that the machine cannot recognize. Delegation of the evaluation functions to the human, and of normative aggregation to the machine also appears to be in agreement with the needs of the operators. Often, subjects will ignore or reject computer aids that assume too great a level of task responsibility. However, Ferguson and Jones (1969) found that subjects readily accepted roles as utility estimators. Their study involved the computer presentation of many dimensions of possible results and the operators assigned weights to the various dimensions.

3.3 Model Learning and Parameter Estimation

3.3.1 Pattern Classification. The adaptive, goal-directed portion of ADDAM is resident in the utility estimator and the utility-based strategy recommendations. The dynamic utility estimation technique is based on the principle of a trainable multi-category pattern classifier. A multi-category pattern classifier receives patterns of data and responds with a decision to classify each of the patterns in one of R categories (Nilsson, 1965). 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

$$g_i(\overline{X}) = \overline{W}_i \cdot \overline{X} \quad \text{for } i = 1,2, \ldots, R \qquad (3.1)$$

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

$$g_i(\overline{X}) > g_j(\overline{X}) \tag{3.2}$$

for all j = 1,2, ..., R; i ≠ j.

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

$$\overline{W}_i' = \overline{W}_i - d \cdot \overline{X} \qquad \text{(Punish)} \tag{3.3}$$

$$\overline{W}_k' = \overline{W}_i + d \cdot \overline{X} \qquad \text{(Reward)} \tag{3.4}$$

where d is the correction increment.

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 can be terminated after some finite number of steps.

3.3.2 The Dynamic Utility Estimator. Probably the easiest way to describe the utility estimation procedure is by example. For instance, a key application of ADDAM is in intelligence gathering. Here, the dynamic utility estimator, shown schematically in Figure 3-5, classifies pattern vectors

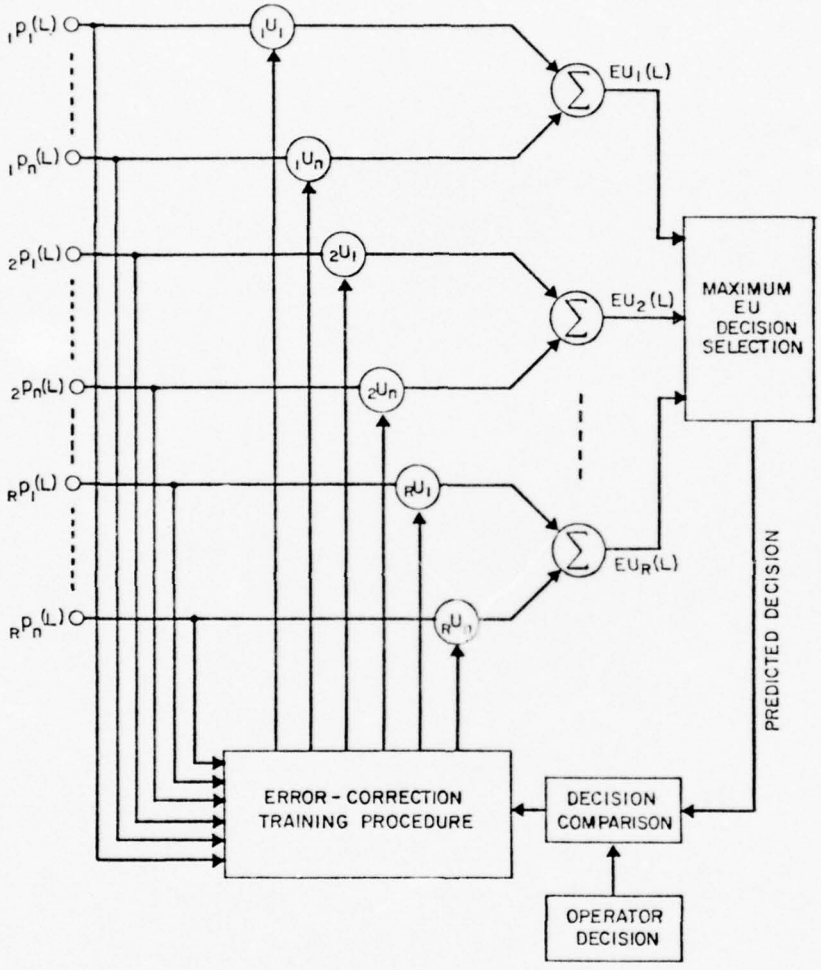


FIGURE 3-5. SCHEMATIC REPRESENTATION OF DYNAMIC UTILITY ESTIMATOR

$$\bar{P} = [{}_1p_1, {}_1p_2, \cdots {}_kp_i \cdots] \tag{3.5}$$

whose components, ${}_kp_i$, are a function of the probability that an object of type i is present and the reliability k of the sensor. These probabilities are weighted by the corresponding utilities ${}_kU_i$ of object sensing. The discriminant functions are then the expected utilities of each sensor decision. The utility estimator computes the EU of each sensor at each location on the board and selects that sensor at each location (including a "null" sensor) 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. A more detailed explanation of the learning algorithm and its underlying assumptions may be found in Davis, Weisbrod, Freedy and Weltman (1975).

3.4 System Characteristics

3.4.1 Closed Loop Features. It is evident from the structural description and from the system behavior that the operator modeling and action recommendation portions of ADDAM comprise a closed loop, adaptive system. The basic closed loop nature of the utility modeling program can be seen in the simplified diagram of Figure 3-6.

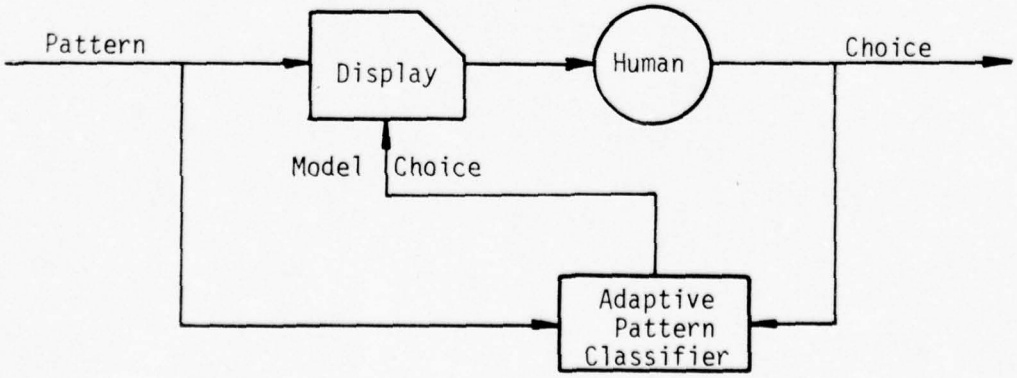


FIGURE 3-6. STRUCTURE OF UTILITY MODELLING PROGRAM

Closed loop behavior is present since the modeling system compares the feedback of the system output (the model choice) with the desired output (the operator choice) and uses this error as an input to the controller. More important, the system is adaptive, since the pattern classifier does not simply operate according to a preset function driven by the error signal, but it adjusts its parameters (the utility weights) to minimize succeeding errors. That is, the behavior of the system is modified toward the goal of maximally predicting operator choice behavior.

3.4.2 Adaptive Features. The adaptivity of the model can be shown empirically in two ways (Gaines, 1972):

(1) Statically: If the operator shows consistent behavior over each of a series of identical cycles of decisions, the system error will be reduced with each cycle.

(2) Dynamically: After training, the system is normally predicting operator behavior satisfactorily. If the operator then changes to a new strategy, prediction will be first unsatisfactory and later satisfactory (this is the definition of a compatibly adaptive system (Gaines, 1972)).

3.4.3 System Performance. Analysis is relatively straightforward when only the system's modeling aspects are considered. If the model predictions are correct, no training takes place; if the predictions are incorrect, the utilities are adjusted. As the DM learns the task and approaches a steady state behavior, the variability of the utility estimates normally approaches a steady state. If the operator behaves most of the time in a manner which is consistent with the model, the variability will be small.

If his behavior is erratic there may be a great deal of variability. Measuring the changes in the utility matrix, therefore, can be used to evaluate the stability of the utilities. The UMD (Utility Matrix Difference) score described in the next section is one such measure.

When both modeling and recommendations based on the modeling are utilized, the picture becomes more complicated. Now the pattern that the operator responds to includes model-based recommendations along with the probability pattern, so that a normative processing of the original input pattern is provided. Since this processing is based on his previously observed behavior, it should lead to greater consistency, speed, and effectiveness in recurrent situations, but it may result in inappropriate predictions for completely new circumstances. These characteristics are typical of predictive displays. The predictions are only accurate if future behavior can be estimated from previous observations. And, as for any estimator, a time lag is present. Thus, with a major structural change in the environment, the recommendations may be based on irrelevant data, and could slow the operator's adjustment. Kunreuther (1969) states that this type of lag can be minimized by including only recent decisions or by exponentially weighting the observations according to their age. The learning program in effect performs such a weighting. A second form of compensation may be accomplished using derivitive feedback of the consistency and rate of movement of the estimated utilities. This is an example of quickening, where the operator is informed of rate of change of parameters, and adjusts his behavior accordingly.

In essence, the modeling system attempts to capture and track the transfer function of the human -- the relationship between the input and output of the decision process -- by operating in parallel with the operator. The estimated transfer function is then used to process the inputs, and the operator's task is changed to one of supervision and correction. Fortunately, formation of this transfer function using the

perceptron-based ADDAM system does not require the explicit knowledge or communication of a utility function. The assumption of the linear model and use of a criterion of maximum prediction is sufficient to capture behavior. In fact, the perceptron does not even require completely accurate observations. Because of its linear form, failures in one or more of its features typically affects performance little (Felson, 1975). In this sense, Felson states that a perceptron-based decision system resembles the human brain: it functions in terms of patterns, and if some elements of the pattern are missing or corrupted with noise, the brain in effect reconstructs them.

3.4.4 Performance Evaluation. The characteristics of the dynamic utility estimator have been evaluated in a variety of decision contexts -- a simulated fishing fleet surveillance task (Davis, Weisbrod, Freedy, and Weltman, 1975), computer aided training (May, Crooks, and Freedy, 1976), and in a simulation of anti-submarine warfare (see Section 4). In all of these applications, the estimates of multiple dynamic utilities typically converged rapidly to stable and distinct values. A representative example of the behavior is given in Figure 3-7. Here the number of utility adjustments per decision is graphed as a function of decision cycles for a specific sensor (this sensor was usually deployed 3 to 4 times per cycle). Initially each decision results in a utility adjustment. Eventually, adjustment is made only about one in ten decisions. 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, during which few adjustments take place and the machine approaches a systematic behavior. In this typical case, major adaptation was completed after about five decision cycles.

Similarly, Figure 3-8 illustrates the adjustments and convergence of the estimated utilities of an expert electronics technician (May, Crooks

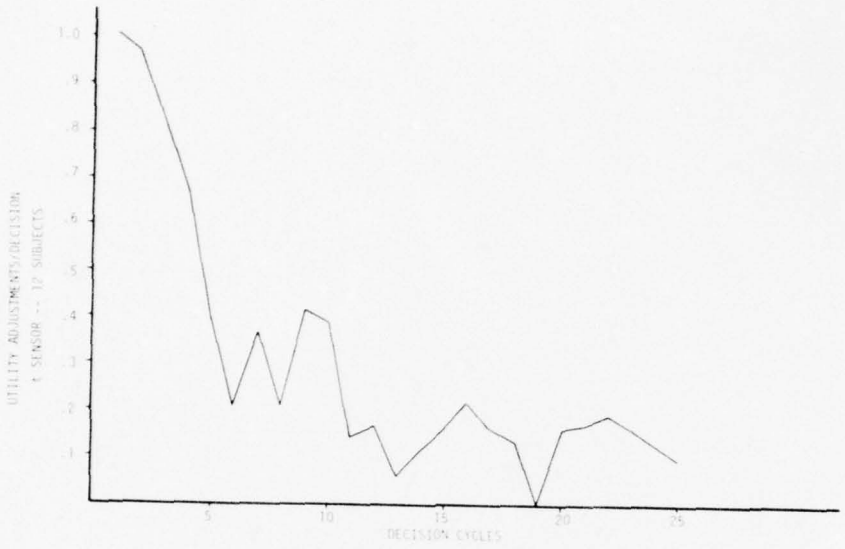


FIGURE 3-7
NUMBER OF UTILITY ADJUSTMENTS PER DECISION
AS A FUNCTION OF DECISION CYCLES

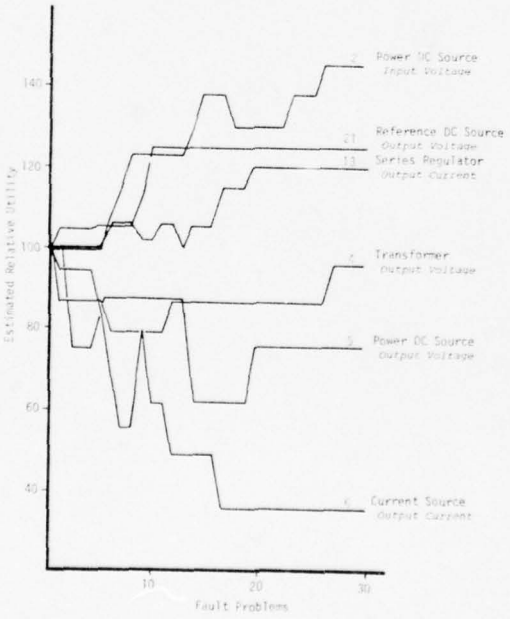


FIGURE 3-8
RELATIVE UTILITY FOR KEY MEASUREMENTS
AS A FUNCTION OF FAULT PROBLEMS
(Expert Technician)

and Freedy, 1976). These utilities represent the apparent preferences for key measurements in a fault diagnosis task. Again, convergence to distinct values was seen.

Of course convergence to stable values is only useful if these values are accurate in predicting operator decisions. ADDAM has been found to be quite effective in prediction, with accuracies ranging from 75% in the computer aided training task (May, et al, 1976) to 95% correct predictions in the simulated intelligence gathering task (Davis, Weisbrod, Freedy, and Weltman, 1975). Additional evidence of applicability was seen in the high correlation (.82, $p<.01$) observed between model estimated and operator expressed preferences in the intelligence task (Weisbrod, Davis, Freedy and Weltman, 1974).

3.5 Decision Measures

One of the advantages of using both descriptive and normative models as the basis of decision aiding is the extensive performance evaluation possible. Two important types of performance measures are directly obtainable. These may be termed (1) decision outcome measures, indices that monitor actual decision effectiveness, and (2) decision quality measures, reflecting the logical soundness of the actions prior to observing the consequences.

3.5.1 Decision Outcome. Outcome measures are, in general, the ones typically used in evaluation of system performance. Their purpose is to define performance in terms of objective, readily-available criteria of cost and achievement. They focus on the actual outcome of an exercise, and include such variables as speed, accuracy, error rate and type, costs, etc.

In the long run, the outcome measures are the true criteria of performance. But while they indicate the actual amount of goal attainment,

these external measures seldom identify the specific deficiencies of behavior. The decision quality measures described in the next section reflect to some degree the decision processes and thus help to identify deficiencies. Nevertheless, the process-oriented quality measures are in the end tied to external performance.

3.5.2 Decision Quality. Quality measures, are used to evaluate the standard of decision making regardless of outcome. This is done through the utilization of normative criteria which classify the decision according to the expected outcomes of the chosen actions. Such performance measures are extremely valuable, because in aiding, one wishes to focus on the decision before its result is determined by the probabilistic outside world (Nickerson and Feehrer, 1975). In many cases, correct decisions have poor consequences, and vice versa. Measures of decision quality include: (1) deviation from maximum expected utility, (2) deviation from maximum expected value, (3) decision consistency, and (4) utility analyses. Each of these is described separately below.

Deviation from Maximum Expected Utility (DEU). The expected utility of a decision is defined as the expected (probability weighted) utility for the possible outcomes of the action selected. If the probabilities and utilities involved are accurate, the DM can do no better than pick the alternative which yields the maximum expected utility. Complete self-consistency seldom occurs, however, and even though a decision maker exhibits a stable preference structure, any randomness of behavior will result in a proportion of decisions which are suboptimal. The extent of the randomness or suboptimality may be measured by the average difference in expected utilities between the optimally derived and the actually chosen alternative:

$$DEU = \frac{1}{N} \sum_{i=1}^{i=N} (EU_{ji} - EU_{imax})$$

where

DEU	=	Deviation from maximum expected utility
N	=	Number of trials
EU_{ji}	=	The DM choice j at the ith trial EU_i
EU_{imax}	=	Max EU at ith decision

Deviation from Maximum Expected Value. The objective payoff gain expected from the placement of a given sensor configuration, $\Delta EV(\underline{S})$, is a measure of the value of the information which can be obtained from this configuration. $\Delta EV(\underline{S})$ is measured in terms of the difference in the expected payoffs resulting from optimal status decisions before and after obtaining the information (Wendt, 1969). Specifically, define:

$\underline{S}$	$\triangleq$	Sensor Deployment Vector
$\underline{D}_i(\underline{S})$	$\triangleq$	The ith possible outcome vector resulting from vector $\underline{S}$
$h_t(j)$	$\triangleq$	Hypothesis that an object of type t is in location j
$P(h_t(j))$	$\triangleq$	*A priori* probability that $h_t(j)$ is true
$P(h_t(j)\|\underline{D}_i(\underline{S}))$	$\triangleq$	*A posteriori* probability that $h_t(j)$ is true given $\underline{D}_i(S)$

For a given outcome vector $\underline{D}_i(\underline{S})$, a rational decision maker should select that hypothesis $h_t(j)$ whose *a posteriori* probability is maximum. Thus,

$$\hat{j}_t(\underline{D}_i(\underline{S})) = \max_j{}^{-1} [P(h_t(j)|\underline{D}_i(\underline{S}))].$$

The expected value of the maximum likelihood status decisions which would be made if outcome vector $\underline{D}_i(\underline{S})$ occurred is

$$EV(\underline{D}_i(\underline{S})) = \sum_t \Big| P(h_t(\hat{j})|\underline{D}_i(\underline{S})) \cdot V(h_t(\hat{j})) - (1-P(h_t(\hat{j})|\underline{D}_i(\underline{S}))) \cdot V(\overline{h_t}(\hat{j})) \Big| ,$$

where

$V(h_t(\hat{j}))$ $\triangleq$ The payoff if $h_t(\hat{j})$ is true

$V(\overline{h_t}(\hat{j}))$ $\triangleq$ The payoff if $h_t(\hat{j})$ is not true

and the summation is taken over all object types.

The expected value gain resulting from a given sensor configuration can thus be computed:

$$\Delta EV(\underline{S}) = \sum_i [P(\underline{D}_i(\underline{S})) \cdot EV(\underline{D}_i(\underline{S}))] - EV(\underline{D}(\underline{S}_o)) - C(\underline{S})$$

Where

$EV(\underline{D}(\underline{S}_o))$ $\triangleq$ The expected value of deploying no sensors

$C(\underline{S})$ $\triangleq$ The cost of deploying $\underline{S}$

Decision Consistency. Decision consistency is a measure of the stability of the Decision Maker's behavioral patterns. Untrained DM's often alter their tactics using inappropriate criteria. Typically, they

modify their behavior on the basis of an inadequate number of preceding trials (Tversky and Kahneman, 1974). Thus, decision consistency is an important measure which defines whether the DM has reached a stable tactical decision strategy. The degree to which the utilities converge provides an easily obtained measure of decision consistency. The level of convergence can be defined by a mean utility matrix difference score (UMD) (Weisbrod, Davis, Freedy, and Weltman, 1974). This measure examines the deviation in estimated utility between successive trials and provides a direct measure of consistency. This measure is computed as follows:

$$UMD\ (T_1,t_2) = \sum_{k,i} |{}_kU_{it_2} - {}_kU_{it_1}|$$

where ${}_kU_{it}$ is the elements of the utility matrix at the t cycle for the kth alternative and the ith outcome. The UMD score is a measure of the variability of the utility values from cycle t_1 to t_2. In the analysis in Chapter 4 a global measure is used, which summarizes the variability of the utilities for the entire session. The session UMD score is the sum of the single-cycle UMD scores from the start of the session, t_o, to the end of the session, t_e. It is defined as:

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

Utility Analysis. The estimated utility function for a particular trainee reflects a value structure which is characteristically predictive of his decision behavior. It is possible to compare this value structure with the value structure of an expert, with organizational values, or with objective criteria of performance and determine deviation from required or optimal values. The utility analysis can also define (1) the relative values

of the DM for various decision outcomes, (2) sensitivity to inter-scenario events, and (3) decision information value. Patterns or clusters of utilities can be also used to characterize specific traits of the decision maker and classify them into profiles of behavior.

Applications. The deviation from maximum expected utility, the utility matrix difference score, and portions of the utility analyses were used in the current series of experiments (Section 4). The deviation from maximum expected value has been tested and validated, and will be used in forthcoming studies.

4. EXPERIMENTAL STUDY OF ASW DECISION MAKING

4.1 Introduction

4.1.1 Overview. Our analysis of adaptive decision aiding and of its potential application to tactical situations led to a variety of questions requiring experimental validation; for example: Does model-based aiding actually contribute to improved information gathering? Are operators able to exploit the aiding given them? Is the aiding primarily effective in improving decision rate or decision quality? These and other questions were the genesis for a full-scale experimental study conducted as the final portion of this research program. The investigation was organized around a realistic ASW task simulation, highly evocative of actual submarine localization and tracking. This chapter reviews briefly our previous experimental work, before presenting descriptions of the current ASW task simulation, a synopsis of the experimental procedures, and a summary and interpretation of the results obtained.

4.1.2 Previous Studies. Three years of analytical and experimental efforts have gone into the development and evaluation of the ADDAM system. Previous to the current work, the experimental investigations were aimed at determining system characteristics, validating the applicability of the approach, and characterizing its sensitivity to operator behavior.

The initial task simulation involved the gathering and evaluation of intelligence about the elements of a fishing fleet as they moved in an expanse of ocean. The operator's functions included efficient placement of sensors of varying capabilities, reliabilities, and costs, in order to maintain surveillance of the fleet. The initial study (Weisbrod, Davis, Freedy, and Weltman, 1974)[1] established the effectiveness of the adaptive decision modeling approach. In this study, the estimates of multiple

[1]An Initial Study in Dynamic Utility Convergence and Decision Aiding.

dynamic utilities converged quickly to stable and distinct values, and the model was found to be very accurate (95%) in predicting each of the six operator's decisions. The adaptive decision model was also found to be sensitive to individual differences in decision strategies. A high correlation (.82, $p<.01$) was seen between sensor preferences expressed by the operator and those determined from the model.

A subsequent study (Davis, Weisbrod, and Freedy, 1975)[2] investigated the effects of model-based aiding on operator behavior. On the average, the operators receiving sensor-placement recommendations performed significantly ($p<.001$) more consistently than the unaided operators. In this analysis, consistency was measured by the deviation from maximum expected utility, essentially the digression from his own calculated optimum (see Section 3.4). Also, the aided subjects completed a significantly greater number of decision cycles during the task, and had less inter-subject variability than the unaided subjects.

4.1.3 Current Approach. The studies to this point had addressed two primary factors: (1) the validity of the adaptive EU model in describing DM choice behavior; and (2) the acceptance of aiding based on the EU model. The emphasis of the current study was to obtain evidence of man-machine system performance improvement under conditions of aiding and performance feedback, to extend the aiding to probability aggregation, and to obtain a greater degree of realism by changing the task to one of ASW coordination.

The complete aiding system under study thus includes the functions of probability aggregation, utility estimation, and strategy recommendation. As applied to the ASW task, probability aggregation involves updating environmental and sensor data when new information is received. Utility estimation attempts to capture the operator's preferences for sensor responses, and recommendations for sensor placement result from the derived utilities. Each of these functions will be described in the coming sections.

---

[2]An Experimental Study of Aiding Effectiveness

A key aspect of the current study was the feedback of system performance, allowing self-evaluation and sharpening of the operator's strategies. Explicit presentation of performance drives the behavior toward organizational goals and provides a criterion for calibrating internal quality measures.

4.2 Simulated ASW Task

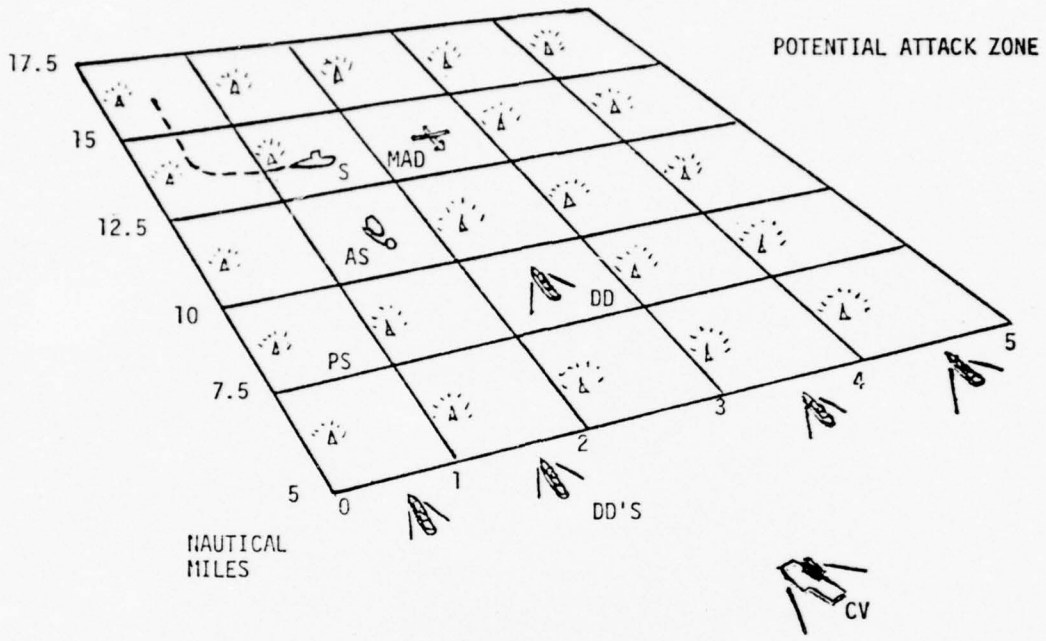
4.2.1 The ASW Environment. The current simulation derives directly from the salient features of actual ASW localization and tracking exercises. The task is an attempt to capture the decision making processes and dynamic flavor of the ASW situation, while retaining some degree of experimental rigor and, for economy's sake, much of the previously developed simulation. It is assumed that the simulated ASW task begins after a hostile submarine has been detected and fixed wing aircraft have deployed sono-buoys over the entire potential attack zone. The attack zone is assumed to remain stationary with respect to the aircraft carrier. Figure 4-1 illustrates the configuration. Several other assumptions, which idealize the behavior of the elements, have also been made. They include the following:

(1) Time and space are both assumed to be discrete. Sensors, submarines, and other objects can be located only to the nearest "sector" of the attack zone. Time advances in discrete steps of 15 minutes.

(2) Each time interval is sufficient for any type of sensor to reach any sector of the attack zone.[3]

---

[3]A forward rate of 10 knots is assumed. Thus in 15 minutes the fleet moves 2.5 nautical miles, or out of a given sector, while aircraft at approximately 200 kts and helicopters at 100 kts are able to reach the furthest sector.



CV AIRCRAFT CARRIER
S HOSTILE SUBMARINE
DD DESTROYER
MAD MAGNETIC ANOMALY DETECTOR (FIXED-WING AIRCRAFT)
AS ACTIVE SONAR (HELICOPTER)
PS PASSIVE SONAR (SONOBUOY)

FIGURE 4.1. CONFIGURATION OF ASW SIMULATION

(3) Objects moving in the attack zone are assumed to move at a speed of no more than one sector per time cycle. Objects can only move north, south, east or west. Objects do not move out of the attack zone.

(4) All sensors are removed at the end of each time cycle and must be re-deployed at the start of the next cycle.

The operator's task is to track the movements of a hostile submarine and, less critically, a non-hostile, mobile object[4] (a whale) as they move in the attack zone. Both move about on an ocean which is represented by a 5 by 5 map grid. The elements move just prior to the start of each task cycle. They can remain stationary or move to an adjacent grid location to the North, South, East, or West. The elements are not able to move off the map, nor to move diagonally.

Neither the operator nor computer aiding system are able to observe the movements of the fleet's elements directly. Their only access to the environment is through sensors which the operator deploys at selected map locations. These sensors differ both in their abilities to detect different types of objects, and in their reliability. A helicopter, for example, can only detect a floating submarine and has high reliability. A destroyer can detect any type of object, also with high reliability; a destroyer, however, is more costly to use. Five different types of sensors are available. Their properties are summarized in Table 4-1. Briefly, the sensors can detect

---

[4]A second, non-hostile object is introduced to divide the evaluator's attention and occasionally (when in close proximity to the submarine) serve as a bogey. This second object is detected (or reported) by some, but not all, of the types of sensors. A whale plausibly fulfills the requirements of this object.

TABLE 4-1. SENSOR PROPERTIES

SENSOR TYPE	LEVEL OF DETECTION	RELIABILITY	COST
Sono-Buoy	Object*	Very Low	Very Low
Mag. Anomaly Detector A	Submarine Floating Submarine Resting	Low	Low
Mag. Anomaly Detector B	Submarine Floating Submarine Resting	Medium	Medium
Helicopter	Submarine Floating	High	Medium
Destroyer	Submarine Floating Submarine Resting Whale	High	High

*A Sono-Buoy reports that an object has been detected but does not identify whether it is a whale or submarine.

objects only at the locations where they are deployed. They are not able to detect objects at adjacent locations. The operator has an unlimited number of each type of sensor, but he is not able to deploy more than one sensor at each grid location. Also, all sensors are removed automatically at the end of each decision cycle and have to be redeployed at the start of the next cycle.

4.2.2 Decision Task. The operator's task is to monitor the movements of the submarine and whale 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. He also receives decision aiding generated by the ADDAM system. This decision task sequence is illustrated in Figure 4.2. The following paragraphs give a more detailed account of what the operator sees and does.

In performing the decision task, the operator sits in front of a graphics display terminal and presented with a simple representation of the environment situation. Figure 4.3 illustrates the graphics display. The operator deploys a pattern of sensors using keyboard inputs and subsequently receives responses from the sensors. 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 4.2 The figure shows that the magnetic anomaly sensor at location C3 has detected a resting submarine. This is represented by the SR in the output column.

The operator reads the sensor outputs, and with the help of sensor output evaluation aiding (described in Section 4.2.3), decides where the objects are. He then enters his status decisions of object location, as illustrated in squares c3 and e5. This status reporting is important, since

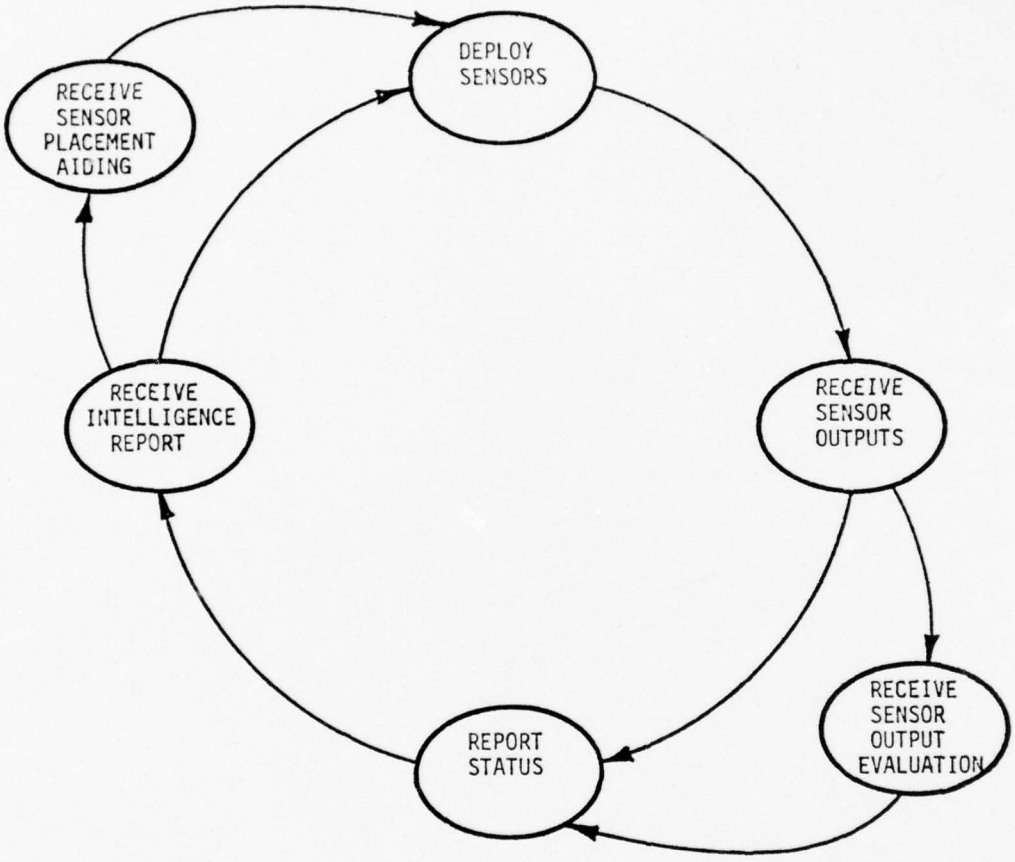


FIGURE 4-2. DECISION TASK CYCLE

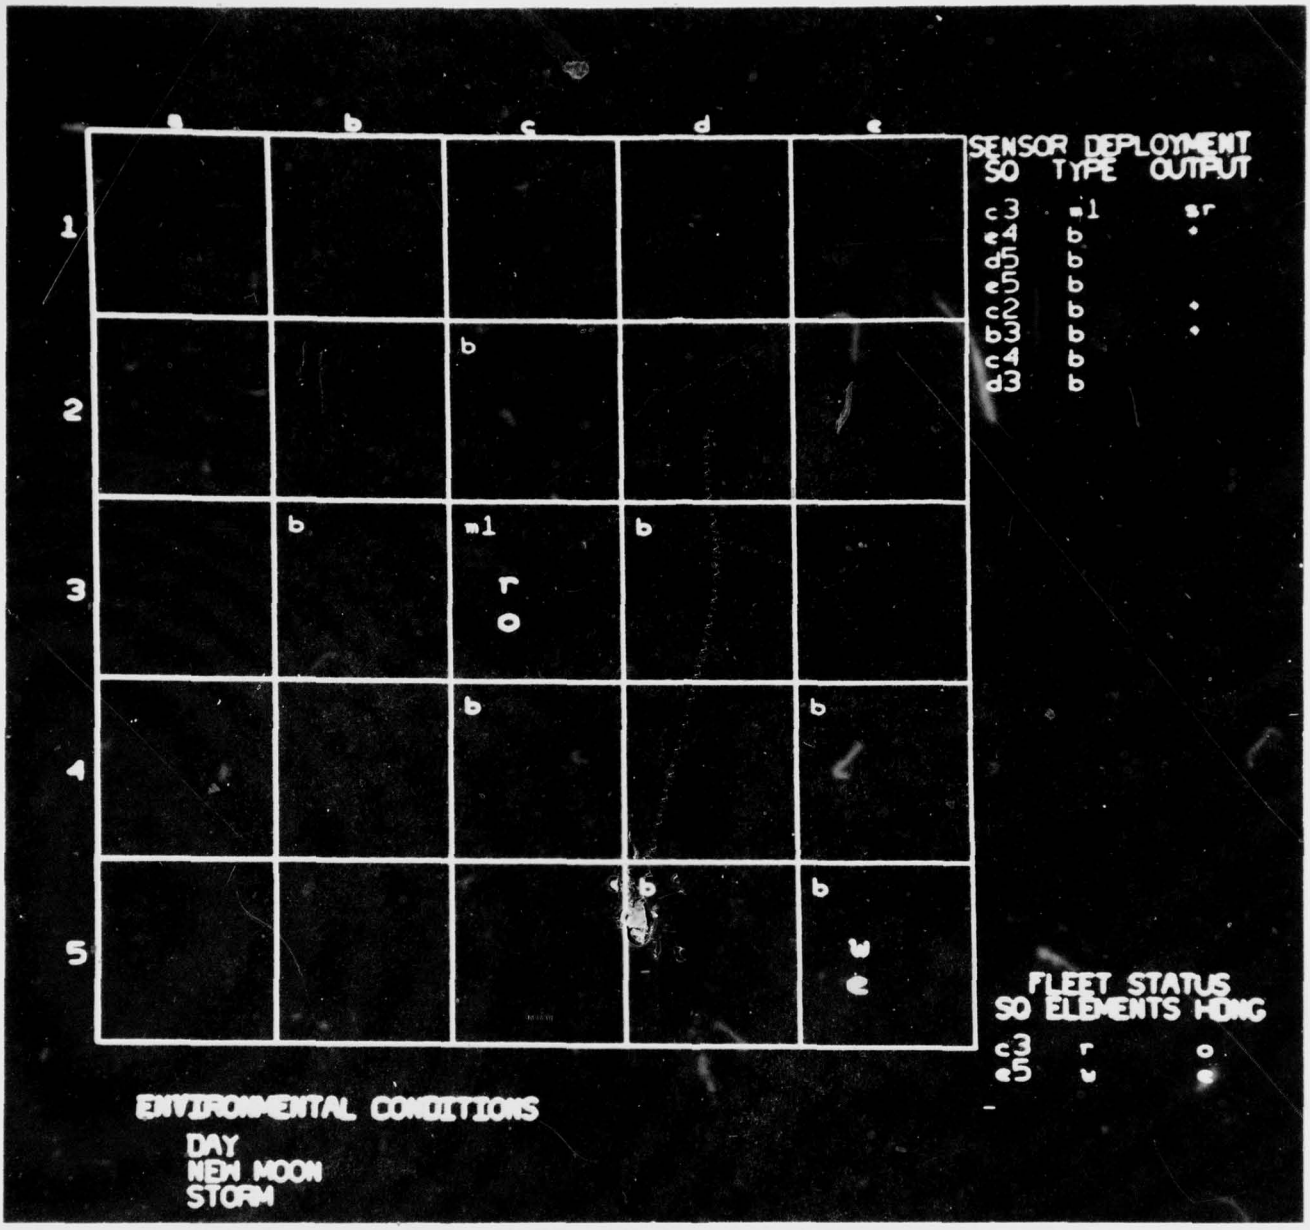


FIGURE 4-3. ASW GRAPHICS DISPLAY

the operator is not only scored on the basis of its accuracy, but the accuracy of the ensuing intelligence report depends on the correctness of the status estimate.

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 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, as shown in the typical intelligence report illustrated in Figure 4-5. After receiving the intelligence report, the operator also receives decision aiding in the form of sensor deployment recommendations. These recommendations (see Section 4.2.3 for description) 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.

Figure 4-4 schematically summarizes the flow of information between the DM and ADDAM. The upper loop represents the main flow of task information. The DM transmits sensor deployment decisions to the real world dimulation and receives sensor outputs. He then makes his status decisions and receives an intelligence report which he uses to make sensor deployment decisions. The organizational information received by the operator includes sensor properties and costs, general strategy guidelines, and general task instructions.

The lower loop in Figure 4-4 represents the decision modeling and aiding loop. The operator's sensor deployment decisions and the intelligence report are automatically input to the decision model. At the experimenter's option, the DM receives decision aiding from the model.

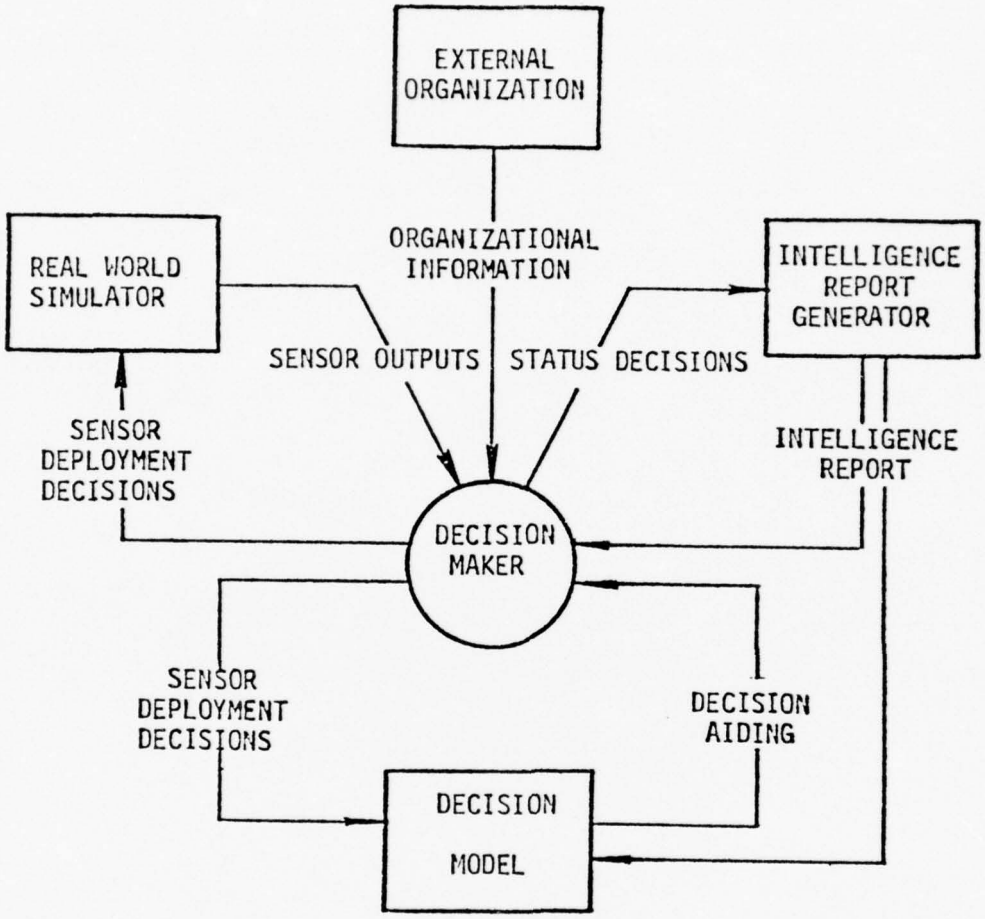


FIGURE 4-4. ADDAM INFORMATION FLOW

4.2.3 Decision Aids. Three types of decision aiding report are generated by ADDAM during the course of a decision cycle. They are discussed here in the order which they occur in the task sequence.

(1) Intelligence Report. This report is derived from the operator's report on the status of the objects being tracked and from expert assessments of the behavior of these types of objects. ADDAM assumes that the evaluator has correctly reported the location and heading of each object and, by aggregating the conditional probabilities of state transformations (elicited from experts), makes a Bayesian estimate of their next location. Thus, the intelligence report contains the probability that an object will be in each sector of the attack zone. Figure 4-5 illustrates a typical intelligence report. The first column lists the sector. The other columns list the probabilities that the submarine (or the submarine on the ocean bottom) and the whale will be in each sector.

(2) Sensor Output Evaluation. This report has the same format as the intelligence report. The probabilities are obtained by using the sensor outputs to update the intelligence report probabilities according to Bayes' Rule. Figure 4-6 shows a typical updated report.

(3) Suggested Sensor Placements. This report is based upon adaptive estimates of the evaluator's utilities for information received from the sensors. It consists of sensor deployment suggestions which maximize the evaluator's expected utility for information gain.

INTELLIGENCE REPORT

	PROBABILITIES		
SECT	SUB	BOT	WH
A1	0	0	48
B1	0	0	25
E1	12	0	0
A2	0	0	25
D2	18	0	0
E2	1	49	0
E3	18	0	0

FIGURE 4-5. INTELLIGENCE ANALYSIS REPORT

SENSOR OUTPUT EVALUATION

SECT	SUB	BOT	WH
A1	0	0	72
B1	0	0	28
E1	9	0	0
E2	0	2	0
E3	9	0	0

FIGURE 4-6. TYPICAL SENSOR OUTPUT EVALUATION

Performance Feedback. The current simulation provides the operator with system performance feedback allowing self-evaluation and sharpening of his decision strategies. Each five decision cycles the following information is presented.

(1) Cost of the sensors deployed
(2) Points for correct detections
(3) Penalties for incorrect status reports
(4) Score, consisting of the points achieved minus the penalties and cost of sensors incurred.

This information is provided for the immediately preceding five decision cycles along with the accumulated total score.

4.3 Methodology

4.3.1 Experimental Hypotheses. The following experimental hypotheses were tested:

(1) The ADDAM system can capture the decision maker's utilities, and predict his sensor preferences, in a simulated ASW task under conditions of system performance feedback.

(2) Decision makers provided with an integrated complement of decision aiding will obtain improved system performance scores.

(3) Aided operators, as individuals and as a group, will be more consistent in their decision behavior.

(4) Aiding will increase operator decision throughput.

4.3.2 Independent Variable. The independent variable under evaluation in the current study was the complement of decison aids. The experimental group was divided into two groups, one of which did not receive aiding.

4.3.3 Dependent Variables. The following measures were used to evaluate the effectiveness of aiding system.

System Performance. The performance measures provided as feedback to the operator during the simulation serve as well as system performance measures. These system performance indices include points, penalties, sensor costs, and accumulated score. The performance score was defined as:

$$\text{Score} = \text{Payoff} - \text{Cost}$$

where

$$\text{Payoff} = \text{Points} - \text{Penalties}$$

Table 4-2 shows the points credited for correct status reports, and penalties deducted for incorrect ones. Cost was the cost of the sensors allocated. Operators were told to maximize their score, and performance feedback was provided every five trials throughout the training and test sessions.

TABLE 4-2. PAYOFF MATRIX FOR STATUS REPORTS

	CORRECT (POINTS)	INCORRECT (PENALTIES)
Submarine	7	-5
Whale	3	-1

Decision Rate. One of the factors contributing to system efficiency is the speed or the number of decisions per unit time. Since all experimental sessions were of the same duration, a reasonable measure is the number of decision cycles completed.

Decision Consistency. The measure of decision consistency is the deviation from maximum expected utility. This measure, defined explicitly in Section 3.5, reflects the extent of correspondence between the normative model and actual behavior.

4.3.4 Experimental Design. A one way experimental design was utilized with a single treatment level across two groups. The experimental group received the sensor recommendation and probability update decision aids. This aided group was compared to a control group which performed the task without benefit of aiding.

4.3.5 Subjects and Procedures. Twelve male subjects were recruited from a nearby Air National Guard Unit. They represented the type of personnel who might interface with computer-aided command systems. The subjects' ages ranged from 20 to 51 years. Six of the subjects were officers ranking from Captain to Lieutenant Colonel and six were non-commissioned officers. The subjects were divided into two groups, an experimental group and a control group, each containing six subjects. Officers and NCO's were divided equally among the two groups.

Each subject performed the task during three sessions of 1-1/2 hours duration. The first two sessions served as training sessions. During the initial session, the subjects were given instructions on system operation and information gain and expected value strategies, and were provided hands-on experience with the equipment for familiarization with input formats. The second session served as a practice session. Here, the

decision maker refined his approach and attempted to maximize his score. The third session was the evaluation session in which the experimental groups were actually compared.

Just prior to the third session, the subjects in the experimental group were introduced to the display of sensor recommendation. They were given instruction on how to accept, reject or modify the aiding, and an explanation of how the sensor recommendations (aiding) were derived. Both groups received instruction on the importance of the final session in determining the bonus they would receive for their performance. The subjects were paid on an hourly basis and were told they would receive a bonus as a function of their score as compared with the scores obtained by the other participants.

The experimental group received probability update aiding during all three sessions, and sensor recommendations during the third evaluation session. The subjects in the control group did not receive aiding at any time, although data was collected on the sensors predicted by the EU model and compared with the sensors actually used.

4.4 Results and Discussion

4.4.1 General. The subjects learned the task procedure readily and by the middle of the second training session could efficiently handle the task requirements. The system responded to the subjects choices rapidly, and by the end of the third training session was able to satisfactorily predict the individual subject's choice behavior.

4.4.2 Results. The preformance scores of the subjects during the test session are shown in Table 4-3. Most importantly, the aided group obtained significantly higher scores than the control group ($p<.05$; t test, 10 df).

TABLE 4-3. MEANS AND STANDARD DEVIATIONS FOR THE TWO GROUPS ON THE PERFORMANCE MEASURES

MEASURE		AIDING	CONTROL	TEST
NUMBER OF DECISION CYCLES	$\bar{X}$	43	37.5	$p<.05$ t test,
	SD	3.1	4.6	10 df
POINTS	$\bar{X}$	290	224.3	$p<.05$ t test,
	SD	28.7	50.5	10 df
PENALTIES	$\bar{X}$	63.7	77.8	
	SD	8.7	18.2	
COST OF SENSOR	$\bar{X}$	68.8	61.2	
	SD	8.7	18.2	
SCORE	$\bar{X}$	161.8	86.2	$p<.05$ t test,
	SD	40.2	64.9	10 df
DEU	$\bar{X}$	1307.5	2642.2	$p<.05$ Fmax test,
	SD	772.2	2471.4	5 df

Part of this score improvement seems to be due to an increase in number of decision cycles completed during the test session ($p<.05$) and part appears to be the result of an increase in decision quality. The increase in decision quality is apparent from the increase in number of points received for correct status reports ($p<.05$), while the sensor cost component of the performance score was not significantly different for the two groups (both in terms of total cost and cost per decision cycle). The penalty component of the performance score was slightly higher for the unaided control group but the difference was not statistically reliable.

It is interesting to note that besides obtaining almost twice the score of the control group, the aided group displayed considerably less variability between subjects. This intersubject consistency is evident from the difference in standard deviation (40.2 aided versus 64.9 control) and appears even more pronounced when the coefficients of variation are compared. The coefficient of variation is 100 times the standard deviation divided by the mean, its purpose is to normalize for the size of the mean. The calculated values are 24.9 for the aided group and 132.8 for the control group.

4.4.3 Decision Consistency. The individual performance consistency for both groups, in terms of deviation from maximum expected utility per trial, was calculated as discussed in Section 3.5. It is seen from the data in Table 4-3 that as a group, the aided subjects demonstrated a high degree of decision consistency, while the control group showed greater extremes of behavior (F_{max} test, homogenity of variance, $p<.05$).

4.4.4 Attitudes and Acceptance. Responses to a questionnaire administered after the final experimental session tend to confirm the performance findings. Five of the six aided subjects stated that aiding speeded up the decision processes, and only one of the six felt that aiding hindered his potential performance. Three of the aided subjects stated that aiding

definitely improved performance by increasing confidence, consistency, and cost-effectiveness. Finally, all subjects experiencing aiding reported the task as easy to master, while half of the six control subjects described the task as difficult.

4.4.5 Discussion. This study confirms the earlier findings that aiding improves decision speed and quality. The time improvement was demonstrated under operator-paced conditions with no externally imposed time constraints. It is to be expected that even greater improvements could be demonstrated under conditions of time stress, where the environment may change independent of the decision maker's timing. The increase in decision speed with aiding was earlier hypothesized to be the result of shifting the decisions to simpler cognitive processes. The non-aided DM must confront each decision using involved processes of problem structuring, recall and evaluation, whereas the aided operator may depend largely on recognition and refining of the logically derived machine recommendations.

The major finding of this final study is that aided subjects obtained higher system performance scores. Examination of the score components revealed that the aided subjects obtained a slightly higher number of points per trial, and about 40% fewer penalties per trial compared to the control group. This, combined with the fact that aided subjects as a group were about 15% faster, resulted in a significant 87% increase in performance score.

The improvement of performance in the aided group seems due in large part to facilitation of consistency of decision making. When a decision maker is consistent in his strategy, he can obtain sufficient data from a stochastic environment with which to evaluate the soundness of his approach. A decision maker who overacts to outcomes and changes his approach before a representative sample of performance is obtained may find it difficult to

hit upon a viable strategy. Conversely, a decision maker who is rigidly consistent may adopt a suboptimal approach and fail to modify it appropriately. It is suggested taht aiding avoids these two extremes by facilitating the development of appropriate decision behavior. Some evidence that this may be the case is found in the highly variable behavior of the unaided subjects compared to the aided group.

4.4.6 Implications. Further analysis of the experimental data and extention to plausible real world situations yielded some important insights to the potential benefit of the decision aiding system. An immediate consideration is the wide range of consequences in operational systems. The system performance score, of course, is a direct function of the selective importance of detection and allowable tolerance of false alarms. In fact, if we assume that incorrect reports should be penalized proportionately more than correct reports are rewarded (a not unreasonable assumption in ASW, where false reports are highly dangerous), then the performance improvement ratio becomes greater as a function of the relative importance of the two factors. For illustration, Table 4-4 shows the potential improvement ratio of aided/unaided scores as a function of the importance weight for penalties. This adjustment is made according to the following definition of the score:

$$\text{Score} = [\text{Points} - W(\text{Penalties})] - \text{Cost}$$

At the equal importance weight of 1 the improvement ratio is 1.88; at a value of 1.5, the ratio is 2.70; at a value of 2.0, the ratio rises to 12.50, and higher values would give much larger ratios. If it is possible to extend the present results to high risk situations, the potential benefits may be considerable.

TABLE 4-4. SCORE AND IMPROVEMENT RATIO AS A FUNCTION OF THE IMPORTANCE WEIGHTS FOR PENALTIES.

W	SCORE = [Points - W (Penalties)] - Cost		
	UNAIDED	AIDED	IMPROVEMENT RATIO
1.00	86.2	161.8	1.88
1.25	65.8	141.6	2.15
1.50	46.4	125.6	2.70
1.75	26.9	109.7	4.07
2.00	7.5	93.8	12.50

Additional impact is foreseen in situations of high complexity and time stress. Previous studies have shown that operators resort to various types of stress-reducing strategies in the face of rapid response time requirements (Tversky and Kahneman, 1972; Wright, 1975). Added complexity in the face of time stress further reduces decision quality (Hayes, 1962). ADDAM's provision of logically derived decision recommendations unloads the operator of much of the burdensome aggregation functions of decision making. Major improvements in the quality of decision making under risk, time stress and complexity are thus expected.

4.4.7 Value of Information. The ADDAM system recommends sensor placements based on the expected utility (EU) of the associated information to the individual operator. Thus among the unique on-line measures available from the system is the EU value of each recommended sensor placement; this is how much each particular piece of information contributes to the EU available from all of the information recommended to the operator (as specified by his model) on a particular cycle.

Figure 4-7 graphs the buildup in percent of available EU as the contributions of decreasingly important sensors are added to the information pool. That is, Sensor No. 1 contributed the greatest incremental EU, Sensor No. 2 the next greatest, and so on. Figure 4-7 was constructed using a number of cycles in which eight sensors were recommended for a single subject. There was little difference in the curve over cycles, as evidenced by the relatively small range of percentage values at each sensor rank. Evaluation of other subjects yielded similar results.

The important aspect of this curve is its shape, which shows that sensors of decreasing importance add less and less EU to the total. In other words, the curve indicates that eliminating the lower utility information would have relatively small effect on the total expected utility. For example,

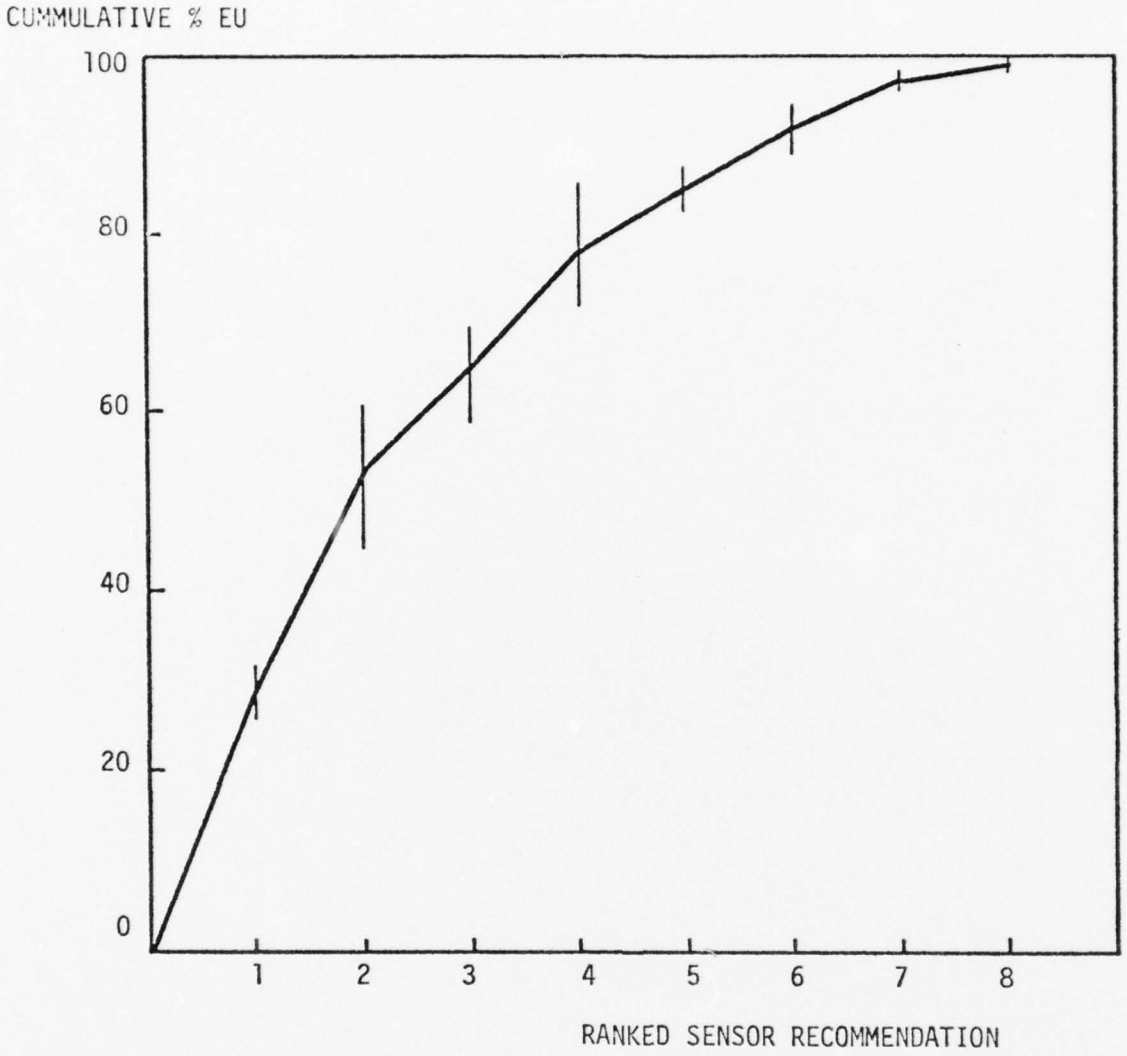


FIGURE 4-7

THE CUMMULATIVE CONTRIBUTION TO TOTAL
EU OF ALL SENSORS IN A RECOMMENDED SET
(Selected Sample)

if the operator were restricted to the best five sensor returns, or 63% of the previous information, his total EU would remain at 85% of the maximum; a more drastic reduction to the three best sensors, 38% of the previous information, still retains 65% of the available EU.

This selected analysis, which looks at the relative contribution of sensors to the total expected utility, suggests that operators tend to expend resources to obtain too much information. This is particularly notable from the quantity of the low quality sources requested.

Even in this optimum case, i.e., considering information with only net positive EU to the operator, the gain from the last ranked sensor is only about 10% of the gain from the first ranked sensor, yet each requires the same attention and response. These results suggest that, in this case, filtering the low-ranked information would have little effect on either acceptance (directly related to EU) or on performance (in this well-practiced and relatively-easy task, utilities probably reflect accurately the value of the information to task performance). In the general case, where the operator is deluged with information of zero or negative utility as well, filtering on a utility basis should have a significant positive effect on overall performance. (This observation is confirmed by the ASW study reviewed in Section 2.3.3.)

4.5 Conclusions

The experimental results showed that the integrated complement of aiding provided by ADDAM was effective in improving system performance. The 88% improvement in score on the ASW task simulation was partially attributable to a small but significant increase in the number of decision trials completed during the session. But most of it appeared due to the better overall quality of the aided decisions. That is, the aided operators incurred slightly higher

costs, but received a much greater return in points, and a lower number of penalties. Decision consistency, as measured by mean deviation from maximum expected utility, was significantly enhanced for the aided group, as in previous studies. And also in replication of previous studies, the improved performance of the aided group was accompanied by decreased intra-group variability.

Almost as important as the results was the power the ADDAM system lent to analysis. The normative expected utility framework of ADDAM allowed examination of the internal quality of the operator's information acquisition and action decisions. Measures of internal consistency, biases, incremental utility of data, and other derivations of the EU model were seen to have a variety of evaluation and training uses. Among the more trenchant are redirection of suboptimal behavior by feedback of inconsistency, and that of control of presentation of information according to the expected incremental utility.

The potential of the ADDAM approach in high risk, time stressed, or complex decision situations also appears strong. The model-based aiding inherent in the system not only unburdens the operator of difficult recall and aggregation tasks, but it helps to insure acceptance by incorporating the operator's own values in its recommendations. While additional experimental verification is necessary to sustain these statements, the present data lends support to the idea that model-based aiding will have its greatest impact in high difficulty, high risk situations.

5. TACTICAL APPLICATIONS OF ADAPTIVE DECISION AIDING

5.1 Introduction

This chapter describes the potential role of adaptive decision aiding in the support and improvement of military tactical decision making. It begins by describing the nature of the tactical decision making task from a users' viewpoint in terms of (1) the decision task itself, (2) the environment in which tactical decision must be made, (3) certain characteristics of tactical decision makers, and (4) the interactions of these elements. The chapter then outlines the basic requirements for good tactical decision making and points out a number of problems adversely affecting a tactical decision maker's ability to make good decisions. These problems are outlined and illustrated by one, rather detailed, example of a current tactical decision making situation -- anti-submarine warfare (ASW). The question of tactical decision aiding is approached by mapping the general discussion of tactical decision making into the decision theory paradigm, and describing how adaptive decision aiding not only supports, but also improves decision making capability. The chapter concludes with a number of short illustrative examples of tactical decision making situations and suggests how adaptive decision aiding could be applied in each.

5.2 The Tactical Decision Task

5.2.1 Definition. Military decision making can be divided into two main categories: (1) strategic and (2) tactical. Strategic decision making tends to be global in outlook and concentrates on long range preparations for possible eventualities, establishing policies, and determining in advance what kind of general guidance and mission statements should be given to tactical decision makers in order to aid them in their local, tactical situations. Strategic decision making is not addressed here.

Tactical decision making, on the other hand, refers to the moment-by-moment direction of local situations that occur in the process of carrying out assigned missions. These decisions are characterized by very short time horizons -- a few hours to a few days -- and there is little, if any, time for deliberation. Prior training, previous experience, and standard operating procedures developed in part by strategic planners provide whatever assistance is available to the tactical decision maker.

Tactical decision making can be thought of as an information feedback system, where the moment-by-moment state of the environment leads to a decision that results in action which in turn affects the moment-by-moment environment and thereby influences subsequent tactical decisions. The tactical decision maker typically taps many sources of information, both reliable and unreliable, timely and delayed, reviews the momentary state of the mission, looks for trends, solicits advice from his staff, and then making use of his training, military experience, and general operating guidelines, makes seat-of-the-pants decisions, acting and reacting as the situation unfolds over time.

5.2.2 The Tactical Decision Cycle. Tactical decision making condenses to the process of converting information and values into actions. Figure 5-1 illustrates this cycle. The initial pressure for action comes from the tactical mission objectives, which originate in the strategic plans. If comparison of the moment-by-moment mission events and the desired objectives shows a discrepancy, then the DM must consider alternative actions. These actions typically involve acquisition of additional information or selection of tactical responses. To make this choice logically, the decision maker must review his list of potential actions and consider the possible consequences of each alternative. He then selects the best course of action based on his own values and the information available. This decision is communicated and converted into action by his tactical forces. As the

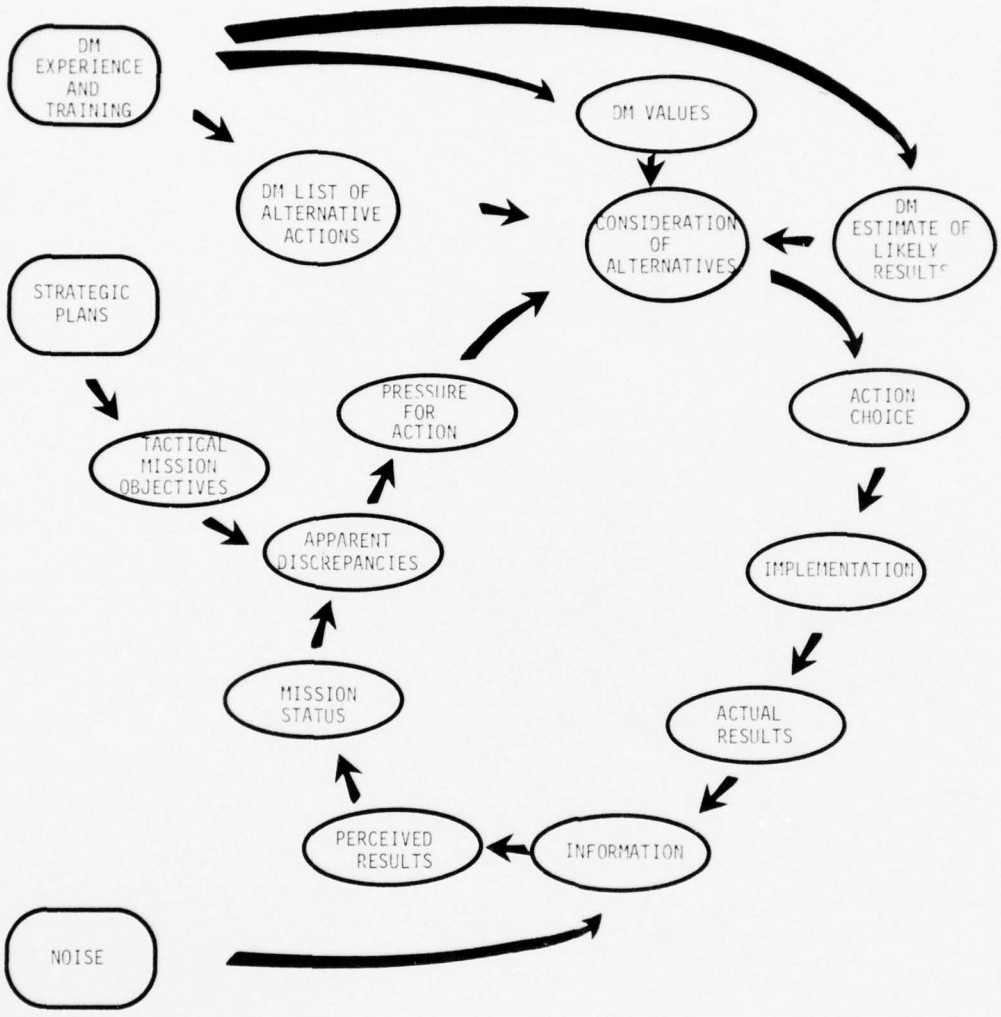


FIGURE 5-1. TACTICAL DECISION MAKING CYCLE

COPY AVAILABLE TO DDC DOES NOT PERMIT FULLY LEGIBLE PRODUCTION

decision is implemented and resources are allocated, perceived results associated with the decisions are observed and reported back to the decision maker. This may consist only of information about the effectiveness of his current decision or may include other, by-product data. The decision maker then compares the perceived results with desired results and again notices any apparent discrepancies. Discrepancies lead to pressure for new action and the cycle repeats itself until the tactical situation is resolved.

The tactical cycle is by no means fixed, as evolution of the conflict will often cause the decision maker to modify his values and probability estimates about likely results. For example, early in a mission a forward lookout may be more valuable than 10 tanks, yet as the situation unfolds, those 10 tanks may be highly preferred to one extra rifleman. Similarly in anti-submarine warfare, passive sonar search is probably more valuable in the early stages of searching than active sonar, yet in the later stages active sonar is typically preferred. In addition to such dynamics, the time constraints, and the information overload, there are also stresses associated with changes in threat status. Often, multiple conflicts develop, as new threats are ascertained. This heightened complexity may occur late in an operation, when fatigue, frustration, hunger and other debilitating conditions have set in. Also, as the operation develops, considerations of the viability of earlier acquired information becomes important.

If tactical decision making is the process of converting information into action, then it is clear that tactical decision making success depends primarily on *what information is chosen* and *how the conversion is executed*. The difference between a good tactical decision maker and a poor one lies at this point. Every tactical decision maker has available a large number of information sources. Each selects and uses only a small fraction of the available information and, even then, makes incomplete and erratic use of that information. The tactical decision maker sets the stage for his

decision (and its results) by his choice of which information to use and which to ignore. He must then decide how to weight the information, each piece of data relative to the other, and then decide how to use this information. How quickly or slowly is it converted to action? How is it converted to action? There are two major factors that influence and constrain tactical decision makers as they chose information and convert it to action: the tactical decision making environment and the tactical decision maker himself.

5.2.3 The Tactical Decision Making Environment. The tactical decision making environment has two major influences on a tactical decision maker: (1) the sequence of moment-by-moment events are unfolding with ever increasing speed, and (2) advances in information systems and communications systems make it possible to drown a decision maker in great quantities of information about a mission's status -- only some of it relevant to decision making. In addition to time constraints and information overload, there are also issues associated with the quality of information. How current or old is it? How important is it in comparison with the other information available? What new information is needed and what old information can be discarded? All are typical of the information related questions that a tactical decision maker must ask; and the response to each is in large part influenced by the tactical environment.

5.2.4 The Tactical Decision Maker. Conversion of information to action is greatly influenced by each individual decision maker's style and experience. Variations in training, knowledge, estimation, abilities, and preferences all shape the operator's interaction with the tactical situation. Thus, each tactical decision maker produces unique sequences of tactical decisions that reflect his own strengths and weaknesses. Moreover, although it is customary to speak of *the* tactical decision maker, this is not necessarily correct in all instances. Over the time span of one tactical situation,

there may be several representatives of the tactical commander, each assuming responsibility in sequence as duty officer. Of course, each brings with him his own abilities or lack thereof in converting information to action. The variations in value systems, experience levels, training, and capabilities in selecting and processing information result in different degrees of smoothness in transitioning from one tactical decision maker to another during a tactical mission.

In addition, individuals can become tired, angry, emotional, forgetful, or simply blunder as time passes. Values for selecting and weighting information, and values for establishing the worth or alternative actions can change during the ebb and flow of a tactical mission. It is during these periods, when preferences for particular actions and results are changing, that prepackaged standard operating procedures are of the least help to actual decision makers. Cookbooks are by their nature static, not dynamic, and cannot possibly allow for all possible mission situations. "We threw away the book" is heard often enough to suggest that rulebooks are not always useful. Mission decision makers must make their own decisions. Each must draw on his strengths (and weaknesses) in selecting and converting information into good tactical decisions, and each is virtually alone.

5.3 A Tactical Example: Anti-Submarine Warfare

5.3.1 Task Description. Anti-Submarine Warfare (ASW) both illustrates the problems discussed above, and affords a rich environment for decision analysis and support. Tactical ASW missions consist of five major functions: intelligence, detection, localization, tracking, and destruction. Intelligence serves to identify the capabilities of classes of submarines and to locate them within a particular ocean basin. Intelligence capabilities include information about the number of submarines in each class, the number and range of missiles, number and type of torpedos, speed and endurance, noise level, and sonar and radar capabilities. Tactical intelligence includes

information about the number of submarines out of port, the number in each ocean basin, operating tactics, special vulnerabilities, and so on.

In the detection stage, an expanse of ocean is searched to see if anything is there. The objective of detection is to make initial contact with hostile submarines. A wide range of sensors are used to "look", including the following: radar, ECM, radio-frequency direction finders, visual sighting systems, magnetic anomaly detectors, passive acoustic methods such as fixed hydrophone arrays, sono-buoy fields, and surface vessel or submarine mounted passive sonar, and active acoustic methods such as active sonar mounted on surface vessels, submarines, helicopters, and ocean floor sensors.

In the localization stage, an object is assumed to have been detected and the objective is to pinpoint its location and identify what it is. The sensors available for this stage are the same as those used for detection, but because the search is much more directed, the patterns of usage and usefullness of the various sensors will differ. The distinction between tactical intelligence on the one hand, and detection and localization on the other, becomes blurred when a submarine's location can be limited to a particular area. The tracking and destruction stages of ASW missions are optional. In peacetime, it is highly unlikely that a submarine will be destroyed once it has been located and identified. Instead, the submarine will probably be tracked to see if it can be forced to surface or until it no longer constitutes a potential threat. In wartime, however, it is far more likely that an attempt will be made to destroy the submarine.

5.3.2 The ASW Decision Environment. A ship-based ASW mission is centered around the Combat Information Center (CIC). In general, the ASW mission is to protect a highly valued vessel, such as an aircraft carrier, from enemy submarine attack. The carrier, either an attack (CVA) or ASW (CVS) carrier, steams along behind a screen of destroyers. An ASW exercise is directed by the screen commander, who is located on one of the destroyers in the CIC and is accompanied by a radar talker, various plotters (i.e., keepers of status-boards, logs, etc.), a sonar talker, and an evaluator.

The evaluator has at his disposal a number of different kinds of sensors, in limited amounts, such as sonobuoys, active sonar, radar, ECM, magnetic anomaly detectors, etc., and a number of means by which these can be deployed, including fixed-wing aircraft, helicopters, and the destroyers themselves. He decides how to deploy the sensors in each stage of the ASW mission, receives and aggregates all of the resulting incoming information, evaluates it, and decides how to continue deployment or how to redeploy his sensors.

Activity is centered around two plotting boards, a long range DRT (Dead Reckoning Tracer) and a short range DRT upon which the courses of the contacts and all fleet elements are plotted. The radar talker receives information and calls out ranges and bearings. This information is plotted on either the long or short range DRT. The sonar talker also shouts out his information and it is similarly plotted. Initially, the long range DRT is used, but when events begin to "speed up" and become critical as the contacts move closer, the plotting activities are transferred to the short range DRT. The evaluator moves back and forth between the two DRT's and keeps the screen commander informed about what is going on. He also directs the air controller who relays information to and from the aircraft pilots of any airplanes or helicopters that happen to be deployed. An ASW exercise can be almost hopelessly complicated by the arrival of additional submarines.

5.3.3 The ASW Decision Cycle. Since the tactical ASW mission is to protect a highly valued vessel by detecting, localizing and tracking enemy submarines, an ASW exercise will typically begin when an ECM device detects a signal, a "suspicious" radar contact is made (called a sinker because it typically lasts for about one sweep of the radar), or when something is detected in the sonar convergence zone. An attempt is made to find out what was detected (attack or missile submarine, a whale, etc.), by beginning the ASW tactical decision making cycle shown in Figure 5-2. The evaluator is the decision maker, while the ASW mission statement and the initial contacts produce the pressure for sensor deployment. Continued repetition of the cycle must redetect the contact, identify, localize, track, and destroy it (if it is wartime and if it is an enemy submarine).

5.3.4 Problems in ASW Decision Making. Simply stated, there is *too much to do* and *too short a time* to do it in for even the most experienced ASW hand in many ASW situations. Information overload occurs as the evaluator is forced to absorb and evaluate excessive amounts of information in an insufficient amount of time. Evaluators, so overloaded, tend to resort to omission, error, approximations, and finally escape. Evaluators have difficulty in estimating likely success of using various sensors. They fail to consider different alternatives or as many as they could (and would) if time permitted. They have difficulty in properly assessing sensor feedback and integrating the information from many diverse sensors. They also exhibit conservatism in revising opinions and biases toward certain information sources to the point that they neglect important and often vital information. They tend to give excessive weight to previously selected decision actions and display an inability to consider several -- let alone all -- feasible alternatives at one time. Their standard tactical procedures,

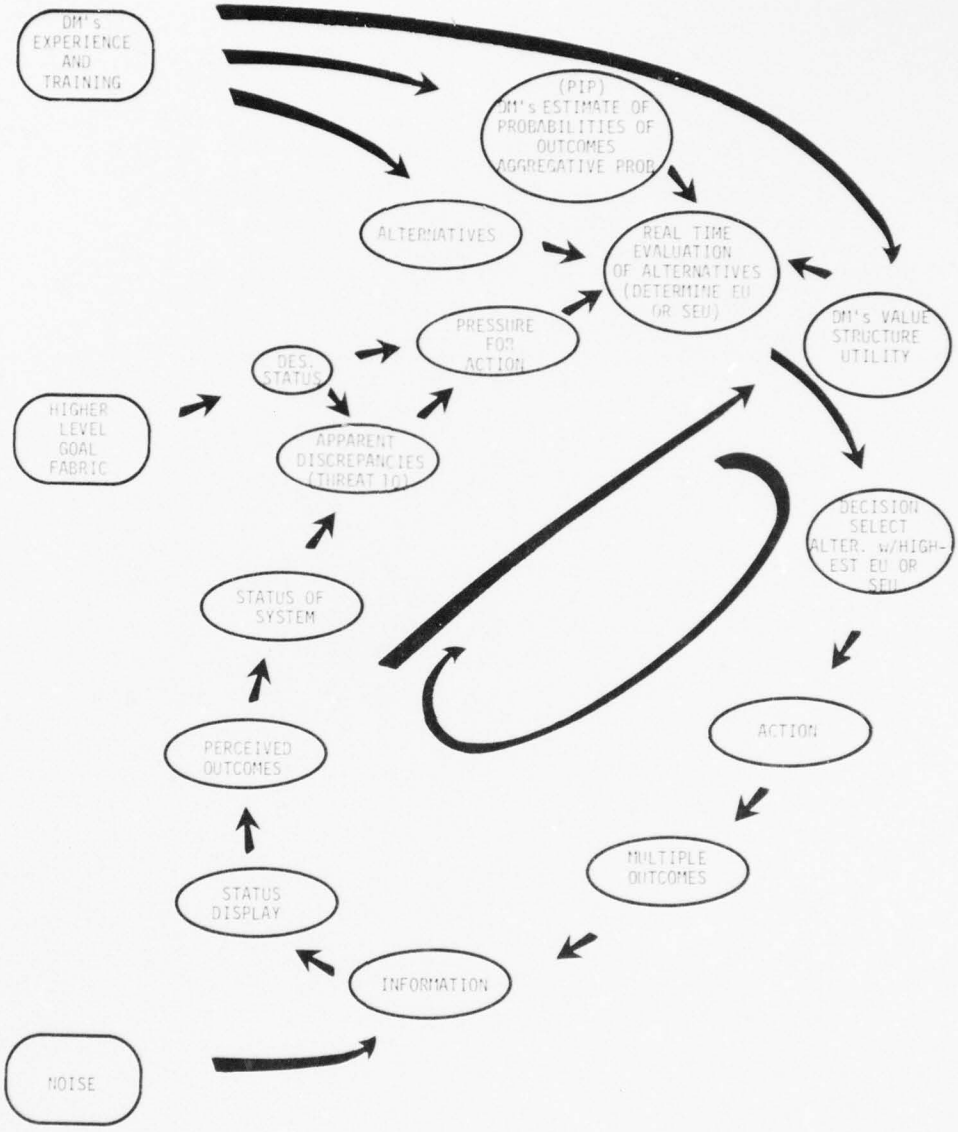


FIGURE 5-2. ASW TACTICAL CYCLE

COPY AVAILABLE TO DDC DOES NOT PERMIT FULLY LEGIBLE PRODUCTION

which often serve as an organizing framework, do not cover anywhere near all the moment-by-moment situations they encounter, so they resort to their own experience and, making the best sense they can out of the information they have in the time available, and, marshalling their current physical and physiological resources, they make a decision.

Analysis of the above problems and deficiencies underscores the need for decision aiding during ASW operations, and also during tactical decision making in general. In the ASW case, some decision and evaluation aids are currently available. But they are limited primarily to standard operating procedures, visual aids, and status boards (e.g., the DRT). These aids do not, however, directly address the main decision task: allocation of ASW sensors, selection of deployment patterns, and evaluation of sensor feedback. The greatest need for decision aiding in tactical situations is in the area of helping tactical decision makers select information and convert it into optimal action in the face of increasing amount of information and decreasing decision time. Moreover, any such aiding processes must be dynamic in nature, so that adjustments can be made as necessary in the feedback system that contains the tactical decision maker.

5.4 ASW and Adaptive Decision Aiding

5.4.1 Decision Theoretical Terminology. Before proceeding, we need to translate the tactical decision making cycle described previously into the terms used in decision theory. Figure 5-3 shows the ASW decision making cycle in decision theoretic terms. The primary differences between Figures 5-2 and 5-3 are as follows. Where previously, tactical decision makers were shown as estimating the likely results of some action they were contemplating, Figure 5-3 shows this in decision theoretic terms as estimating and aggregating the probabilities of action outcomes. Also where the previous figure showed values, Figure 5-3 also shows values, but in the formal sense, that is, a measure of a decision maker's utility or

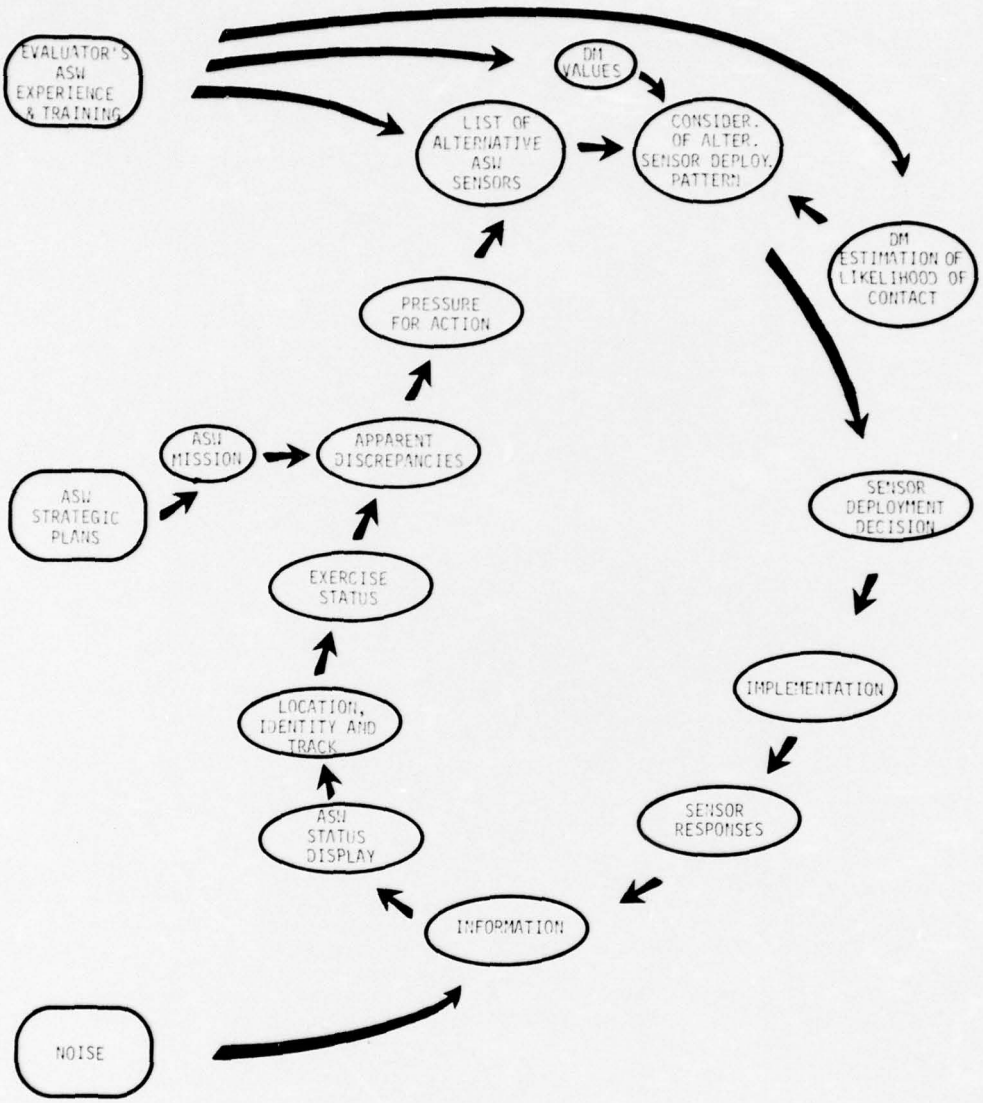


FIGURE 5-3
THE TACTICAL DECISION MAKING CYCLE IN
DECISION THEORETIC TERMS

COPY AVAILABLE TO DDC DOES NOT
PERMIT FULLY LEGIBLE PRODUCTION

preferences for outcomes. Consideration of alternatives becomes evaluation of alternatives by determining the expected utility of each, taking into account its associated probabilities and utilities. Finally, the decision rule is to select the alternative with the highest expected utility. Then, continuing the cycle, as actual outcomes occur, the results are used to revise probability estimates and utility structures.

5.4.2 Tactical Aiding Concept. In ASW the evaluator is the decision maker and his task is to allocate sensor resources in order to locate and track the movements of enemy submarines in some defined portion of an ocean. As the ASW exercise unfolds the adaptive system "watches" the decision to deploy the sensors available in response to the information available. For example, a decision to place a helicopter with its sonar sensor in some location in some situation can have one of two possible outcomes: "positive", indicating the presence of a submarine in the sector, and "negative" indicating the absence of a submarine. (It should be noted here that the sensors are not perfectly reliable in that their responses may be erroneous.) The system provides utility estimates for the various decision outcomes. The utility estimates are a measure of the relative worth of the outcomes of each sensor placement to the ASW evaluator. Thus, the utilities associated with a helicopter are for positive and negative feedback information about the presence of a submarine at locations where the helicopters are deployed. After the evaluator's utilities have been learned, decision aiding is produced for each decision task.

5.4.3 Tactical Aiding Modes. As previously discussed, the capabilities of adaptive aiding derive primarily from its on-line and adaptive character. In the case of ASW, as well as for similar tactical situations, these capabilities lead to aiding potential in a number of key problem areas. Table 5-1 lists the main deficiencies in ASW decision making, and places

TABLE 5-1. ASW DECISION AIDING

	PERFORMANCE DEFICIENCY	DECISION AID
1.	Failure to consider an adequate number of feasible sensor deployment alternatives.	Generate feasible alternatives.
2.	Failure to revise probabilities correctly.	Revise probabilities using Bayes' theorem.
3.	Failure to consider expected utility of each sensor deployment pattern.	Calculate the expected utility of each alternative.
4.	Making suboptimal sensor deployments.	Recommend the optimal sensor deployment to maximize the evaluator's expected utility.
5.	Time constraints. Lack of time for decision making as additional submarines are added.	Provide decision aid at computer speed, using learned utilities.
6.	Shift changes and subsequent loss of tactical continuity.	Continue to provide sensor deployment recommendations; learn new values, if any; give continuity between duty personnel changes.
7.	Human physical and physiological failures	Provide recommendations that are not affected by human factors limitations.

alongside the type of adaptive aiding which could be used to help overcome each. The more important of these -- sensor selection, information diagnosis, and performance feedback -- are discussed separately below.

Sensor Selection and Deployment. This aiding consisting of sensor deployment suggestions would maximize the evaluator's expected utility for information gain, based upon adaptive estimates of the evaluator's utilities for information received from the sensors. The aiding in selection of sensor deployment is obtained from the ability to evaluate rapidly the expected outcomes of plausible decision alternatives in terms of their multiple evaluation criteria, using the utilities it has learned (and continues to learn) from the evaluator. Evaluation criteria which might be considered include the following: (1) expected utility and expected value of each outcome, (2) expected reduction in uncertainty, (3) expected impact on decisions to follow, and so on. It is important to note that in making decisions under uncertainty, it is not possible, *a priori*, to know that the decision will yield the desired result. In such situations one can only make that decision which has the greatest expectation of yielding the desired result.

Sensor Information Diagnosis. Evaluation of sensor data makes up the diagnosis phase of the decision process. The task is to evaluate the outcomes of previous decisions and transform the information gained into a form that can be used to make future decisions. The technology for doing this, called probabilistic information processing, or PIP (Edwards, 1962), is very well established, and is extensively applicable to use in ASW decisions, since probabilities -- and their revision -- plays a major part in any decision. In using PIP, current information about the state of the environment (for example, the location of a submarine) is stored in the form of *a priori* probabilities through Bayes' rule, thus producing *a posteriori* probabilities.

Performance Feedback. The decision biases and inconsistencies of an evaluator can be discovered and reduced by observation and feedback of his behavior. The expected utility model is well suited to this function because of its normative form and feedback of parameters of the model has been shown to redirect behavior to more effective actions (Steeb, et al, 1976).

5.4.4 Conclusions. The foregoing discussion has shown the applicability of an adaptive aiding system to a generalized intelligence gathering task, namely shipboard ASW. The characteristics of this task -- high information load, rapid response requirements, recurrent decisions, and reliance on subjective evaluation -- combine to produce a situation where adaptive techniques are of help. Of course, such situational characteristics are not unique to the operation of shipboard ASW, and in fact are common to a variety of command and control applications. The next section outlines several areas having immediate potential for aiding.

5.5 Illustrative Examples of Tactical Adaptive Decision Aiding

The following examples are discussed briefly to indicate how ADDAM could provide decision aiding across a range of tactical situations.

P-3 Aircraft. This is essentially airborne rather than shipboard ASW where the decision maker is called the TACCO (tactical coordinator). Since the TACCO is in an aircraft and must use various sensors to detect, identify, localize, and track submarines, decision time is critical. Delays in making sensor deployment decisions for as short a time as 30 seconds can cause a lost contact.

Air Traffic Control. In air traffic control situations aircraft arrive and depart from air base runways under the control of air traffic controller. As aircraft volume increases and controllers become

fatigued, the number of incidents (near misses) typically increases. Aiding in action recommendation and alerting to critical situations is needed.

Anti-Air Warfare (AAW). In AAW, time is even more critical than in ASW. Rather than submarines moving at speed around 20-30 knots, incoming aircraft missiles traveling at a speed of 200 knots or more may be threatening a fleet. The decision time is reduced to the point where an entire AAW exercise may begin and end in a very short time frame. Such a time stress necessitates the use of immediate response recommendations, possibly using decision rules developed from previous interaction.

Command and Control. In many tactical situations continuity of command decision making is a problem because, as a tactical situation evolves, duty personnel including the tactical decision maker may change shifts. This happens frequently in high level headquarters where the command is represented by a duty officer position. An individualized system can provide continuity in the face of such duty personnel changes, as well as aid in the information acquisition and decision-making functions.

6. TECHNOLOGY TRANSFER

The experience gained using the ADDAM system has initiated similar developments in a variety of areas. Computer-based aids in such disparate fields as decision training, data routing and pacing, allocation of man/machine function, tactical resource dispatching and others have been realized or are now in development using concepts derived from ADDAM. These projects, in spite of having widely varying goals, share a number of key characteristics:

(1) The operator is confronted with a series of difficult but reasonably well defined decisions. That is, performance is normally sub-optimal and the decision structure -- the possible actions and outcomes -- can be delineated.

(2) Major aspects of the decisions are recurrent. Thus, estimation of behavioral parameters for modeling is possible.

(3) The task performance relies heavily on the subjective evaluations of the operator, rather than on some completely specified objective criteria. Thus operator preferences or utilities play an important part in system direction, and inference of utilities is required.

Beyond these characteristics, the systems are quite general. Ideally, the decisions under consideration may be risky or riskless, single stage or sequential, static or dynamic, information centered or action centered, forced-paced or self-paced, and one-dimensional or multi-attribute. In essence, the ADDAM system has evolved from a specific technique to a potentially generalizable methodology. Current realizations of this evolution are the research and applications programs diagrammed in Figure 6-1 and described in the following paragraphs.

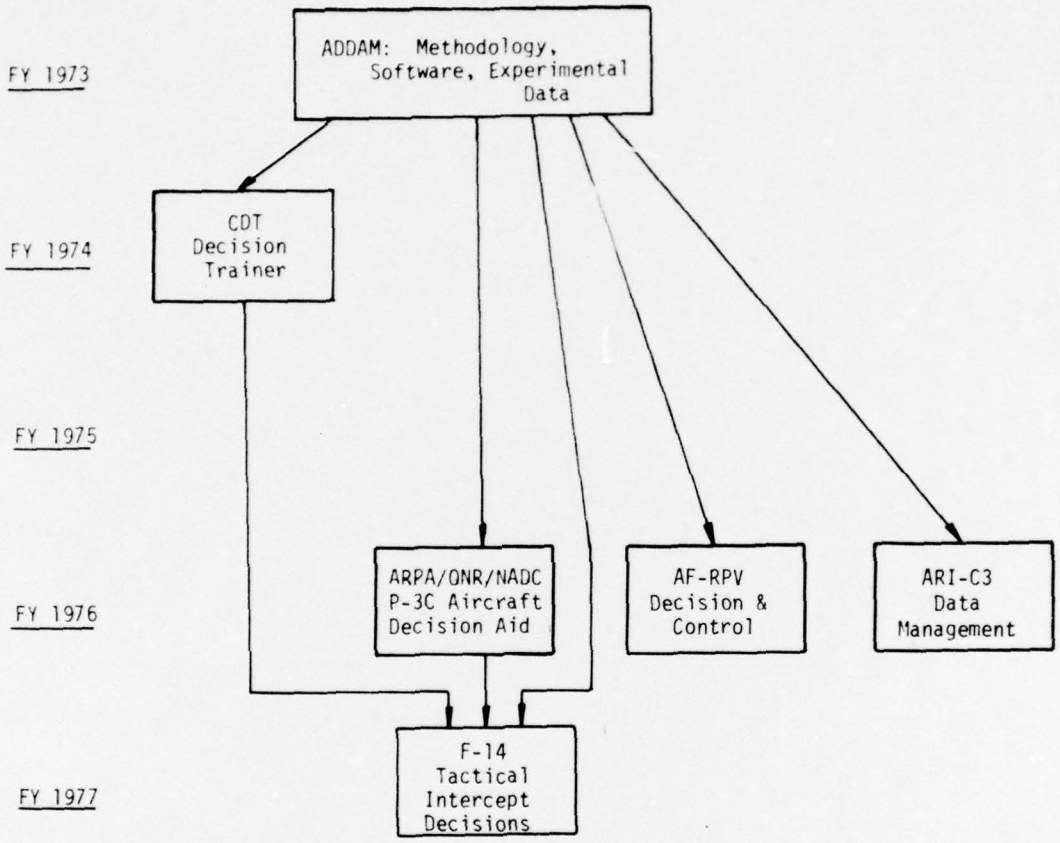


FIGURE 6-1. EVOLVEMENT OF ADDAM TECHNOLOGY TRANSFER

COPY AVAILABLE TO DDC DOES NOT PERMIT FULLY LEGIBLE PRODUCTION

Computerized Decision Training (CDT). The CDT is a direct application of the ADDAM system to training higher order skills of diagnosis and value judgment in decision making tasks. The system incorporates an adaptive computer program which learns the student's diagnosis and decision strategy, compares this strategy to an expert's standard of performance, and adapts the instructional sequence to move the student's performance in the desired direction. The adaptation is accomplished by providing inputs to the heuristic algorithms which generate the instructions and adjust the problem presentation sequence. The instructor model also generates suggested actions in response to student requests for assistance. A prototype system of this type has been shown to effectively train operators in the sequential decisions associated with electronic maintenance (May, Crooks, and Freedy, 1976).

Remote Piloted Vehicle (RPV) Information Management. A multi-attribute variant of the ADDAM system is being developed to select and transmit costly and sensitive down-link information. This is a research program directed toward modeling decision information needs in continuous control of advanced aircraft. Adaptive, autonomous machine decision-making rather than operator aiding is emphasized.

Anti-Submarine Warfare (ASW) Resource Allocation. The current ADDAM application is being modified to aid the specific decisions of the tactical coordination officer in the P-3C ASW aircraft. The ADDAM system will assist in the tasks of target detection and identification by improving the management of sensor and weapon systems. The program may lead to the installation of the decision aid in the P-3C aircraft.

Command and Control (C3) Data Flow. Multi-attribute models of information needs and dynamics are being developed for C3 systems. These models will be used to request, route, and pace information between an interconnected net of human data sources and recipients. The models are

essentially adaptations of the ADDAM structure to the form of the information pacing and addressing matrices of Roby (1968).

F-14 Tactical Intercept Decisions. The decision aiding methodology is being applied to the judgmental tasks associated with tactical intercept operations in the F-14 cockpit. A multi-attribute utility model is being developed to aggregate a number of objective and subjective dimensions to select a course of action. The major dimensions will be elicited from experts using protocol and Delphi techniques, while the dimensions weightings will be estimated in step through simulations.

The ADDAM system is especially suited to capturing decision strategies in dynamic environments. Dynamic decision making is a multi-stage process involving the consideration of a sequence of decisions, in which the results of earlier decisions affect the structure of later decisions (Edwards, 1962). Attempts to make analytical statements of the dynamic processes using extensive trees (Brown, et al, 1974) and dynamic programming (Rapoport and Wallsten, 1972) often meet with intractable complexities or require unrealistic assumptions. Instead, a descriptive approach is used in ADDAM to contend with dynamic situations. That is, the value of an outcome of an action is assumed to be influenced by the changes in the environment and the changes in the availability of future information resulting from the action. In short, the outcome utilities estimated on the adaptive program reflect the operator's conception of both present and future consequences. The model is parsimonious. Also like the brain, only those structural aspects necessary to capture and analyze the situation are included.

7. REFERENCES

Ackoff, R.L. Management Misinformation Systems. Management Science, 1967, 14 : B147-B156.

Albright, A.F. Processing and Display of Information in the National Military Command Center (NMCC). U.S. Army War College (Carlisle Barracks, Pennsylvania) AD#A-012614, 23 May 1975.

Barron, F.H. An Information Processing Methodology for Inquiring into Decision Processes. In White, D.J. and Bowe, K.C. (Eds), The Role and Effectiveness of Theories of Decision in Practice. London: Hodder and Stoughton, 1975.

Beach, B.H. Expert Judgment about Uncertainty: Bayesian Decision Making in Realistic Settings. Organizational Behavior and Human Performance, 1975, 14: 10-59.

Bowman, E.H. Consistency and Optimality in Managerial Decision Making. Management Science, 1963, 9: 310-321.

Brown, R.V., Hoblitzel, C.M., Peterson, C.R., and Ulvila, J.W. Decision Analysis as an Element in an Operational Decision Aiding System. Decisions and Designs, Inc. (McLean, Virginia), Technical Report 74-2, September 1974.

Coombs, C.H. and Pruitt, D.G. Components of Risk in Decision Making: Probability and Variance Preferences. Journal of Experimental Psychology, 1960, 60:265-277.

Davis, K.B., Weisbrod, R.L., Freedy, A., and Weltman, G. Adaptive Computer Aiding in Dynamic Decision Processes: An Experimental Study of Aiding Effectiveness. Perceptronics, Inc. (Woodland Hills, CA), Technical Report PTR-1016-75-5, May 1975.

Dawes, R.M. A Case Study of Graduate Admissions: Application of Three Principles of Human Decision Making. American Psychologist, 1971, 26:180-188.

Dawes, R.M. and Corrigan, B. Linear Models in Decision Making. Psychological Bulletin, 1974, 81:95-106.

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

Edwards, W. Man as Transducer for Probabilities in Bayesian Command and Control Systems. In Shelly, M.W. and Bryan, G.L. (eds), Human Judgment and Optimality, New York: Wiley, 1964.

Felson, J. Artificial Intelligence Techniques Applied to Reduction of Uncertainty in Decision Analysis Through Learning. Operational Research Quarterly, 1975, 26(3):581-598.

Ferguson, R.L. and Jones, C.H. A Computer Aided Decision System. Management Science, 1969, 15(10):B-550-B-561.

Fishburn, P.C. Methods of Estimating Additive Utilities. Management Science, March 1967, 13(7):435-453.

Fischer, G.W. Multi-Dimensional Value Assessment for Decision Making. Engineering Psychology Lab., University of Michigan, Tech. Report 037230-2-T, 1972.

Freedy, A., Hull, F.C., Lucaccini, L.F., and Lyman, J. A Computer-Based Learning System for Remote Manipulator Control, IEEE Trans. on S.M.C., 1971, SMC-1: 356-363.

Freedy, A., Steeb, R., and Weltman, G. Man/Machine Interaction in Adaptive Computer Aided Control: Initial Study. Perceptronics, Inc. (Woodland Hills, CA), Technical Report No. 72-10, November 1972.

Freedy, A., Weisbrod, R., and Weltman, G. Self-Optimization of Task Allocation in Shared Man/Computer Control. Proc. of the IEEE Conference on Decision and Control, Paper #FP5-3, San Diego, California, December 5-7, 1973.

Gaines, R. Axioms for Adaptive Behavior, Int. Journal of Man-Machine Studies, 1972, 4:169-199.

Gilstad, D.W. and Fu, K.S. A Two-Dimensional Pattern Recognizing Adaptive Model of the Human Operator, Proc. Sixth Annual Conf. Manual Control, April 1970.

Green, P.E., Halbert, M.H., and Minas, J.S. An Experiment in Information Buying. Journal of Advertising Research, 1964, 4: 17-23.

Goldberg, L.R. Man versus Model of Man: A Rationale, Plus Some Evidence, for a Method of Improving on Clinical Judgments. Psychological Bulletin, 1970, 73: 422-432.

Goodman, B., Saltzman, M., Edwards, W., and Krantz, D.M. When SEU is Irrelevant. University of Michigan, Department of Psychology, 1971.

Halpin, S.M., Johnson, E.M., and Thornberry, J.A. Cognitive Reliability in Manned Systems. IEEE Transactions on Reliability, R-22:165-170.

Halpin, S.M., Thornberry, J.A., and Streufert, S. The Credibility of Computer Estimates in a Simple Decision Making Task. Purdue University, Tech. Report No. 5, ONR Contract N00014-67-A-0030, January 1973.

Hanes, R.M. and Gebhard, J.W. The Computer's Role in Command Decision. U.S. Naval Institute Proceedings, 1966, 92:61-68.

Hayes, J.R.M. Human Data Processing Limits in Decision Making. In E. Bennett (Ed.), Information System Science and Engineering. Proceedings of the First Congress on the Information Systems Sciences, New York: McGraw-Hill, 1964.

Howell, W. Some Principles for the Design of Decision Systems: A Review of Six Years of Research on a Command/Control System Simulation, AMRL-TR-67-136, Aerospace Medical Research Laboratories, Wright-Patterson Air Force Base, Ohio, September 1967.

Johnson, E.M., Cavanagh, R.C., Spooner, R.L., and Samet, M.G. Utilization of Reliability Measures in Bayesian Inference: Models and Human Performance. IEEE Transactions on Reliability, 1973, R-22:176-183.

Johnson, E.M. and Halpin, S.M. Preliminary Evaluation of a Multi-Stage Bayesian Inference System. The Proceedings of International Conference on Cybernetics and Society, IEEE SMC Soc. October 1972.

Kaplan, R.J. and Newman, J.R. Studies in Probabilistic Information Processing. IEEE Transaction on Human Factors in Electronics, 1966, HFE-7: 49-63.

Kanarick, A.F., Huntington, J., and Peterson, R.C. Multi-Source Information Acquisition with Optional Stopping. Human Factors, 1969, 11: 379-386.

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

Kneppreth, N.P., Gustafson, D.H., Leifer, R.P., and Johnson, E.M. Techniques For Assessment of Worth. U.S. Army Research Institute for the Behavioral and Social Sciences, Technical Paper No. 254, 1974.

Krantz, D.M., Luce, R.D., Suppes, P. and Tversky, A. Foundations of Measurement: Additive and Polynomial Representations, I, New York: Academic Press, 1971.

Kunreuther, H. Extensions of Bowman's Theory on Managerial Decision Making. Management Science, April 1969, 15(8): B-415-B-432.

Lee, W. Decision Theory and Human Behavior. New York: Wiley, 1971.

Levine, J.M. and Samet, M.G. Information Seeking with Multiple Sources of Conflicting and Unreliable Information. Human Factors, 1973, 15(4): 407-419.

Levine, J.M., Samet, M.G., and Brahlek, R.E. Information Seeking with Limitations on Available Information and Resources. Human Factors, 1975, 17(5):502-513.

Levit, R.A., Heaton, B.J., and Alden, D.G. Decision Support in a Simulated Tactical Operations System (SIMTOS). Honeywell, Inc., Contract No. DAHC19-75-C-0008, November 1975.

Levit, R.A., Alden, D.G., Erickson, J.M., and Heaton, B.J. Development and Application of a Decision Aid for Tactical Control of Battlefield Operations. Vol. 1: A Conceptual Structure for Decision Support in Tactical Operations Systems. Honeywell, Inc., Contract No. DAHC19-73-C-0069, August 1974.

Lichtenstein, S., Slovic, P., and Zink, D. Effect of Instruction in Expected Value on Optimality of Gambling Decisions. J. of Exper. Psychol. 1969, 79: 236-240.

Lichtenstein, S. and Slovic, P. Reversals of Preferences Between Bids and Choices in Gambling Decisions. Journal of Experimental Psychology, 1971, 89:46-55.

Luce, R.D. and Suppes, P. Preference, Utility, and Subjective Probability. In Luce, R.D., Bush, R.R., and Galanter, E. (Eds) Handbook of Mathematical Psychology. New York: Wiley, 1965, pp. 42-49.

May, D.M., Crooks, W.H., and Freedy, A. Application of Adaptive Decision Aiding Systems to Computer Aided Instruction: Experimental Studies. Perceptronics, Inc. (Woodland Hills, CA), Technical Report PATR-1023-76-1, January 1976.

Myers, A., Gibb, C.A. and McConville, C. Game Interactions with Humans and Machines and their Relation to Tactical Behavior and Success. Educational Testing Service: 2959(00) RB-63-14, 1963.

Miller, L.W., Kaplan, R.J., and Edwards, W. JUDGE: A Value-Judgment-Based Tactical Command System. Organizational Behavior and Human Performance, 1967, 2: 329-374.

Miller, L.W., Kaplan, R.J., and Edwards, W. JUDGE: A Laboratory Evaluation. Organizational Behavior and Human Performance, 1969, 4: 97-111.

Newman, J.R. Assessing the Reliability and Validity of Multi-Attribute Utility Procedures: An Application of the Theory of Generalizability. SSRI Res. Report 75-7, Univ. of Southern California (Los Angeles), 1975.

Nickerson, R.S. and Feehrer, C.E. Decision Making and Training: A Review of Theoretical and Empirical Studies of Decision Making and their Implications for the Training of Decision Makers. Bolt, Beranek, and Newman (Cambridge, Mass.) Tech. Report NAVTRAEQUIPCEN 73-C-0128-1, July 1975.

Nilsson, N.J. Learning Machines. San Francisco: McGraw Hill, 1965.

O'Connor, M.F., Peterson, C.R., and Palmer, T.J. Stakes and Probabilities in Information Purchase. Organizational Behavior and Human Performance, 1972, 7:43-52.

Peterson, C.R., Schneider, R.J., and Miller, A.J. Sample Size and the Revision of Subjective Probabilities. Journal of Experimental Psychology, 1965, 69: 522-527.

Phillips, L.D. and Edwards, W. Conservatism in a Simple Probability Inference Task. Journal of Experimental Psychology, 1966, 72: 346-357.

Pitz, G.F. and Reinhold, H. Payoff Effects in Sequential Decision Making. Journal of Experimental Psychology, 1968, 77:249-257.

Rapoport, A., Lissitz, R.W., and McAllister, H.A. Search Behavior With and Without Optional Stopping. Organizational Behavior and Human Performance, 1972, 7:1-17.

Rapoport, A. and Wallsten, T.S. Individual Decision Making. Annual Review of Psychology, 1972, 23: 131-176.

Roby, T.B. Small Group Performance. Chicago: Rand McNally, 1968.

Rouse, W.B. Cognitive Sources of Suboptimal Human Prediction, Ph.D. Dissertation MIT, Report DSR 70283-19, September 1972.

Samet, M.G. Checker Confidence Statements as Affected by Performance of Initial Image Interpreter. U.S. Army Research Institute for the Behavioral and Social Sciences (Arlington, VA), Technical Research Note 214, 1969.

Schaffer, L.H. Problem Solving on a Stochastic Process. Ergonomics, 1965, 8: 181-192.

Schrenk, L.P. Aiding the Decision Maker -- A Decision Process Model, Ergonomics, 1969, 12: 543-557.

Schroeder, R.G. and Benbasat, I. An Experimental Evaluation of the Relationship of Uncertainty in the Environment to Information Used by Decision Makers. Decision Sciences, 1975, 6(3): 556-567.

Schum, D.A. Behavioral Decision Theory and Man/Machine Systems. In DeGreene, K.B. (Ed). Systems Psychology. McGraw Hill: New York, 1970.

Slagle, J.R. Artificial Intelligence: The Heuristic Programming Approach. New York: McGraw Hill, 1971.

Slovic, P. and Lichtenstein, S. Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment. Organizational Behavior and Human Performance, 1971, 6: 649-744.

Slovic, P., Fischoff, B., and Lichtenstein, S. Behavioral Decision Theory, Annual Review of Psychology, 28, in press.

Snapper, K.J. and Peterson, C.R. Information Seeking and Data Diagnosticity. Journal of Experimental Psychology, 1971, 87: 429-433.

Steeb, R., Artof, M., Crooks, W.H., Freedy, A., and Weltman, G. Man/Machine Interaction in Adaptive Computer Aided Control: Final Report. Perceptronics, Inc. (Woodland Hills, CA), Technical Report PATR-1008-75-12/31, December, 1975.

Sterling, T.D. Humanizing Computerized Information Systems. Science, 1975, 190: 1168-1172.

Tversky, A., and Kahneman, D. Judgment Under Uncertainty: Heuristics and Biases. Science, 1974, 185, 1124-1131.

Weisbrod, R.L., Davis, K.B., 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, Contract N00014-73-C-0286/NR196-128, ARPA Order No. 2347, Code 455, Office of Naval Research, November 29, 1974.

Weisbrod, R.L., Davis, K.B., and Freedy, A. Adaptive Utility Assessment in Dynamic Decision Processes: An Experimental Evaluation of Decision Aiding. Proc. International Conference on Systems, Man, and Cybernetics, San Francisco, California, September, 1975.

Wendt, D. Utility and Risk. Acta Psychologica, 1970, 34:214-228.

Wendt, D. 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.

Wiggins, N. and Kohen, E.S. Man versus Model of Man Revisited: The Forecasting of Graduate School Success. Journal of Personality and Social Psychology, 1971, 19: 100-106.

Wright, P. The Harassed Decision Maker: Time Pressures, Distractions, and the Use of Evidence. Journal of Applied Psychology, 1974, 59(5): 555-561.

APPENDIX A. SYSTEM DESCRIPTION

1. System Software

The system software is composed of three parts: the real world model, the adaptive decision and training model, and the human interactive routines. All three parts are inter-related, and form a complete information processing cycle. In addition, appropriate printing and display routines are provided for transmitting information to the experimental subject for his attention and updating.

The real world model consists of: (1) a scenario routine to generate and update the scenario and the real world environment, (2) a sensor system routine to associate the sensor outputs with certain reliabilities, and (3) an intelligence generator to determine the intelligence probabilities. These probabilities are used by the subject as a reference and are treated as guidelines for the calculation of the sensor evaluation report.

The adaptive decision and training model handles the following functions: (1) the recommendation of the type and location of sensors according to their expected utility (EU), (2) the evaluation of the impact of sensor outputs on the intelligence report, and (3) the dynamic training of the adaptive decision model to follow desired operator behavior.

The human interactive routines consisted of: (1) the status generated from the subject's actions and the transmission of this data to the real world model, and (2) the deployment of the sensors on the board by the subject.

2. Major Sub-Programs

Figure A.1 is a logic flow diagram of the software system. The major sub-programs are classified as follows:

(a) Real World Model:

SNARIO -- This routine generates scenarios for the dynamic decision tasks. It implements the aggregation of elemental conditional probabilities according to Bayes' rule and modifies the overall probabilities of state transitions. These transitions result in changes in the operational environment.

GENWSE -- This sensor controller determines, according to the environmental conditions and the deployed sensors, what sensor outputs to produce. Probabilistic sensor reliabilities are obtained by consulting a set of random numbers.

PROBST -- This routine is designed to simulate an intelligence expert who knows how the behavior of the environment and the objects but must rely upon status reports from the subject for data about the current status of the objects. It also provides the subject with the probabilities of possible actions occurring based on the information in the subject's status report.

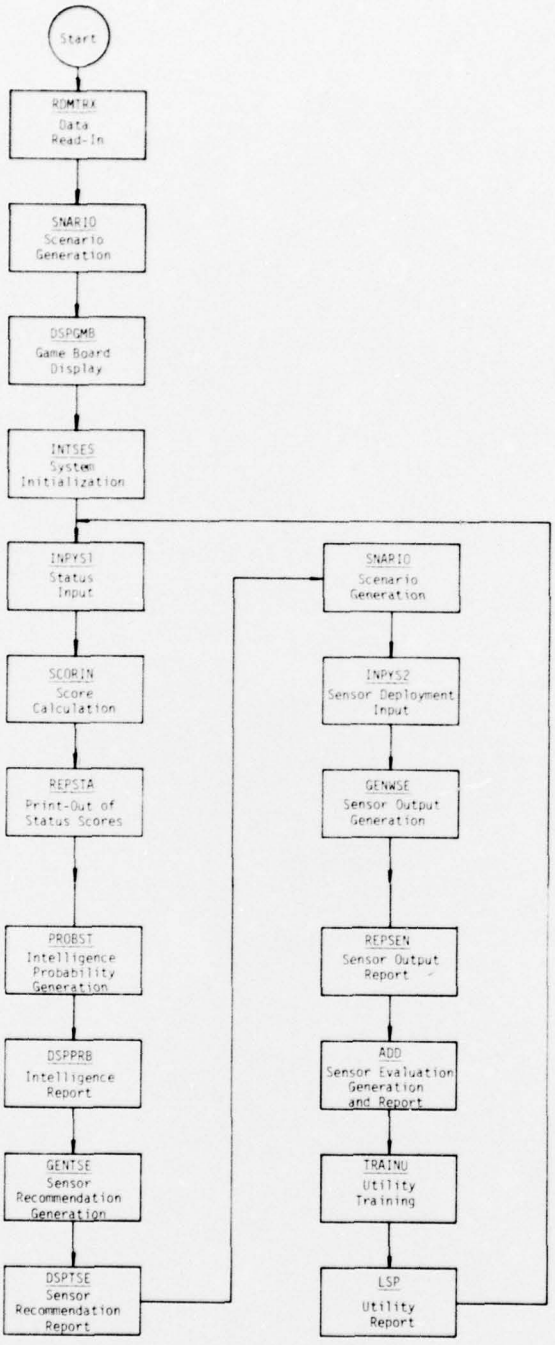


FIGURE A-1. SYSTEM FLOW CHART

(b) Adaptive Decision and Training Model

GENTSE -- The expected utilities for each possible combination of squares are calculated here. Based upon the maximum likelihood algorithm, the sensor with the maximum EU is chosen as the recommended sensor for a given square. The EU of the kth sensor for square ℓ is defined as:

$$EU_{i,\ell} = \sum_{i=1}^{3} [p_{i,\ell} - (1-p_{i,\ell})\beta_i] U_{i,k} (1-p_{i,\ell})(1-\beta_i) \bar{U}_{i,k} - C_k$$

where the symbols are defined as in Section 3.3.2 where i represents the ith type object.

ADD -- Modifications of the intelligence report are made by comparing the current report to the sensor deployment and responses. Training is accomplished by punishing a probability when there is a negative outcome and rewarding when there is a positive outcome; meanwhile, all other probabilities are normalized. The resulting probabilities constitute the sensor evaluation report. The expected values (EV) are also calculated here.

TRAINU -- For each square the deployed sensors are compared to the recommended sensors and the utilities are modified if the predictions are incorrect. Once the utilities are modified, the recommended sensors for all subsequent squares are re-calculated and the new

recommended sensors are used for the subsequent comparisons.

RDMTRX -- Beta errors, sensor costs, initial utilities, and all other performance measurement data are input here.

INTSES -- In this program, all the variables are initialized for the starting of a new experiment.

SCORIN -- Penalties correct actions, scores, and payoffs are calculated here at the end of each cycle.

(c) Human Interactive Routines

INPYS1 -- At the beginning of each cycle, the subject inputs the status based upon the system conditions and his previous experience. These inputs affect subsequent intelligence report probabilities.

INPYS2 -- For localization and tracking, the subject deploys the appropriate sensors. He can also impose the search mode for locating the objects. In search mode, low-reliability sensors are placed over the entire board by a search command.

In addition to the above routines, the display and printout subroutines include: DSPGMB (Display the Game Board), DSPPRB (Intelligence report printout), REPSTA (printouts of status report, and score report), REPSEN (sensor output report printout) DSPTSE (display and printout of sensor recommendations), and UP (printout of utility assessment).

3. Hardware Requirements

The functional relationships of the devices in the simulation task system are illustrated in Figure A.2.

The individual components and their functions are as follows:

1) Computer: An Interdata Model 70 minicomputer system containing the CPU, 48K bytes of core memory and all peripheral I/O interfaces.

2) Graphic Display: An IDI alphanumeric CRT unit display and a keyboard terminal.

3) TTY: A Teletype Model 33 teletypewriter for accepting system commands and printing of decision aiding data.

4) Printer: A Centronics Model 306 line printer for result printout on a hard copy basis.

5) Paper Tape Reader: A Digitronics Model 2540 perforated tape reader for high speed loading of the system into memory.

DISTRIBUTION LIST

Director, Engineering Psychology (5 cys)
Programs, Code 455
Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

Defense Documentation Center (12 cys)
Cameron Station
Alexandria, VA 22314

Lt. Col. Henry L. Taylor, USAF
OAD(E&LS) ODDR&E
Pentagon, Rm 3D129
Washington, D.C. 20301

Capt. Granum
Office of Assistant Secretary of
Defense (Intelligence), Pentagon
Washington, D.C. 20301

Dr. Robert Young
Cybernetics Technology Office
Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Human Factors Plans, OP987P7
Office of the Chief of Naval Operations
Department of the Navy
Washington, D.C. 20350

Personnel Logistics Plans, OP987P10
Office of the Chief of Naval Operations
Department of the Navy
Washington, D.C. 20350

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Assistant Chief for Technology
Code 200
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Office of Naval Research (6 cys)
International Programs
Code 102IP
800 North Quincy Street
Arlington, VA 22217

Aircraft Instrumentation Program
Code 211
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Sensor Systems Program
Code 222
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Fleet Analysis and Support Division
Code 230
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Naval Analysis Programs
Code 431
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Operations Research Program
Code 434
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Statistics and Probability Program
Code 436
Office of Naval Research
800 N. Quincy
Arlington, VA 22217

Information Systems Program
Code 437
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Ocean Technology Program
Code 485
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217

Director, ONR Branch Office
ATTN: Dr. J. Lester
495 Summer Street
Boston, MA 02210

Director, ONR Branch Office
ATTN: Psychologist
536 South Clark Street
Chicago, IL 60605

Director, ONR Branch Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Director, ONR Branch Office
ATTN: Mr. R. Lawson
1030 East Green Street
Pasadena, CA 91106

Dr. M. Bertin
Office of Naval Research
Scientific Liaison Group
American Embassy - Room A-407
APO San Francisco, CA 96503

Dir., Naval Research Laboratory (6 cys)
Technical Information Division
Code 2627
Washington, D.C. 20375

Mr. John Hill
Naval Research Laboratory
Code 5707.40
Washington, D.C. 20375

Dr. Andreas B. Rechnitzer
Office of the Oceanographer
of the Navy
Hoffman Building II
200 Stoval Street
Alexandria, VA 22332

Dr. Heber G. Moore
Hqs., Naval Material Command
Code 0331
Department of the Navy
Washington, D.C. 20360

Mr. Arnold Rubinstein
Naval Material Command, NAVMAT 0344
Department of the Navy
Washington, D.C. 20360

Commander, Naval Air Systems Command
Crew Station Design, AIR 5313
Department of the Navy
Washington, D.C. 20361

Commander, Naval Facilities
Engineering Command
R&D Plans & Programs Division
Code 031A
Alexandria, VA 22332

LCDR T. W. Schropp
Naval Sea Systems Command
NAVSEA 00C3
Washington, D.C. 20362

Commander, Naval Air System Command
Human Factors Programs
AIR 340F
Washington, D.C. 20361

Commander, Naval Air Systems Command
ATTN: Mr. T. Momiyama
Advance Concepts Division, AIR03P34
Washington, D.C. 20361

Naval Electronics Systems Command
Human Factors Engineering Branch
Code 4701
Washington, D.C. 20360

Mr. James Jenkins
Naval Sea Systems Command
Code 06H1-3
Washington, D.C. 20362

Dr. Herbert J. Mueller
Naval Air Systems Command
AIR-310, Research Admin.
Washington, D.C. 20361

Dr. James Curtin
Personnel & Training Analyses Office
Naval Sea Systems Command
NAVSEA 074C1
Washington, D.C. 20362

Bureau of Medicine & Surgery
Operational Psychology Branch
Code 513
Washington, D.C. 20372

CDR Paul Nelson
Naval Medical R&D Command
Code 44
Naval Medical Center
Bethesda, MD 20014

Director, Behavioral Sciences Department
Naval Medical Research Institute
Bethesda, MD 20014

LCDR Thomas Berghage
Behavioral Sciences Department
Naval Medical Research Institute
Bethesda, MD 20014

Dr. George Moeller
Head, Human Factors Engineering Branch
Submarine Medical Research Laboratory
Naval Submarine Base
Groton, CT 06340

CDR James Goodson
Chief, Aerospace Psychology Division
Naval Aerospace Medical Institute
Pensacola, FL 32512

Bureau of Naval Personnel
Special Assistant for Research Liaison
PERS-OR
Washington, D.C. 20370

Dr. Fred Muckler
Manned Systems Design, Code 311
Navy Personnel Research and
Development Center
San Diego, CA 92152

LCDR M. O'Bar
Code 305
Navy Personnel Research and
Development Center
San Diego, CA 92152

LCDR M. Curran
Human Factors Engineering Branch
Crew Systems Department
Naval Air Development Center
Johnsville
Warminster, PA 18974

CDR Robert Kennedy
Human Factors Engineering Branch
Code 5342
Pacific Missile Test Center
Point Mugu, CA 93042

Mr. Ronald A. Erickson
Head, Human Factors Branch
Code 4075
Naval Weapons Center
China Lake, CA 93555

Systems Engineering Test Directorate
Human Factors Section
U.S. Naval Air Test Center
Patuxent River, MD 20670

Mr. Richard Coburn
Head, Human Factors Division
Naval Electronics Laboratory Center
San Diego, CA 92152

Navy Personnel Research and
Development Center
Management Support Department
Code 210
San Diego, CA 92152

Human Engineering Branch
Code 2733
Naval Ship Research & Development
Center, Annapolis Division
Annapolis, MD 21402

Commanding Officer & Director
Naval Coastal Systems Laboratory
Panama City, FL 32401

Mr. John Quirk
Naval Coastal Systems Lab
Code 712
Panama City, FL 32401

Dr. Robert French
Naval Undersea Center
San Diego, CA 92132

Dr. Jerry C. Lamb
Display Branch, Code TD112
Naval Underwater Systems Center
New London, CT 06320

Mr. James L. Long
Weapons Systems Research (N-332)
Naval Education and Training Command
Naval Air Station
Pensacola, FL 32508

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Dr. Alfred F. Smode
Training Analysis & Evaluation Group
Naval Training Equipment Center
Code N-00T
Orlando, FL 32813

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 93940

Mr. J. Barber
Headquarters DA, DAPE-PBR
Washington, D.C. 20546

Dr. Joseph Zeidner
Dir., Organization and Systems
Research Laboratory
U.S. Army Research Institute
1300 Wilson Boulevard
Arlington, VA 22209

Dr. Edgar M. Johnson
Organizations & Systems Research Lab
U.S. Army Research Lab
1300 Wilson Boulevard
Arlington, VA 22209

Technical Director
U.S. Army Human Engineering Labs
Aberdeen Proving Ground
Aberdeen, MD 21005

U.S. Air Force Office of Scientific
Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, D.C. 20332

Chief, Human Engineering Division
Aerospace Medical Research Lab
Wright-Patterson AFB, OH 45433

Dr. Donald A. Topmiller
Chief, Systems Effect. Branch
Human Engineering Division
Wright-Patterson AFB, OH 45433

Lt. Col. Joseph A. Birt
Human Engineering Division
Aerospace Medical Research Laboratory
Wright Patterson AFB, OH 45433

Lt. Col. John Courtright
Headquarters
AMD/RDH
Brooks AFB, Texas 78235

Dr. E. Weiss
Office of Science Information Service
National Science Foundation
1900 Pennsylvania Avenue
Washington, D.C.

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Dr. Stanley Deutsch
Office of Life Sciences
Headquarters, NASA
600 Independence Avenue
Washington, D.C. 20546

Journal Supplement Abstract Service
American Psychological Association
1200 17th Street, N.W. (3 cys)
Washington, D.C. 20036

Dr. Clyde Brictson
Dunlap and Associates, Inc.
115 South Oak Street
Inglewood, CA 90301

Dr. Alphonse Chapanis
The Johns Hopkins University
Department of Psychology
Charles & 34th Streets
Baltimore, MD 21218

Mr. Edward Connelly
OMNEMII, Inc.
410 Pine Street, S.E.
Suite 200
Vienna, VA 22180

Mr. J. G. Curtin
McDonnell Douglas Astronautics
Company - East
St. Louis, MO 63166

Dr. Robert R. Mackie
Human Factors Research, Inc.
Santa Barbara Research Park
6780 Cortona Drive
Goleta, CA 93017

Mr. Alan J. Pesch
Eclectech Associates, Inc.
Post Office Box 179
North Stonington, CT 06359

Mr. Lawrence Shumaker
Woods Hole Oceanographic
Institution
Woods Hole, MA 02543

Dr. A. I. Siegel
Applied Psychological Services
404 East Lancaster Street
Wayne, PA 19087

Dr. W. S. Vaughan
Oceanautics, Inc.
3308 Dodge Park Road
Landover, MD 20785

Dr. Robert Williges
University of Illinois
Institute of Aviation
Savoy, IL 61874

Director, Human Factors Wing
Defense & Civil Institute of
Environmental Medicine
Post Office Box 2000
Downsville, Toronto, Ontario
Canada

Dr. A. D. Baddeley
Director, Applied Psychology Unit
Medical Research Council
15 Chaucer Road
Cambridge, CB2 2EF
England

Dr. David Zaidel
Road Safety Centre Technion City
Haifa
Israel

Prof. Dr. Carl Graf Hoyos
Institute for Psychology
Technical University
8000 Munich
Arcisstr 21
Federal Republic of Germany

ARI Field Unit - Leavenworth
Post Office Box 3122
Fort Leavenworth, Kansas 66027

Dr. Stanley N. Roscoe
Aviation Research Laboratory
University of Illinois
Savoy, Illinois 61874

Dr. William A. McClelland
Human Resources Research Office
300 N. Washington Street
Alexandria, VA 22314

Dr. John Shore
Code 5403
Communication Sciences Division
Naval Research Laboratory
Washington, D.C. 20375

CDR Richard Schlaff
NIPSSA
Hoffman Building I
2461 Eisenhower Avenue
Alexandria, VA 22331

Navy C3 Architecture Division, OP-943
Office of the Chief of Naval Operations
ATTN: LCDR D. A. Spaugy
3801 Nebraska Ave., N.W.
Washington, D.C. 20390

Mr. L. A. Aarons
R&D Plans Division
Office of the Chief of Naval Operations
OP-987C
Washington, D.C. 20350

Commander, Naval Electronics Systems Command
C3 Project Office
PME 108-1
ATTN: LCDR E. Neely
Washington, D.C. 20360

Commander, Naval Electronics Systems Command
Command and Control Division
Code 530
Washington, D.C. 20360

Dr. Marshall A. Narva
Army Research Institute for Behavioral & Social Sciences
Room 236, 1300 Wilson Blvd.
Arlington, VA 22209

Major David Dianich
DSMS
Building 202
Fort Belvior, VA 22060

Mr. William Martin
Analytics Inc.
1405 Colshire Drive
McLean, VA 22101

Dr. Bertram Spector
CACI, Inc. - Federal
1815 N. Fort Myer Drive
Arlington, VA 22209

Mr. Harold Crane
CTEC, Inc.
7777 Leesburg Pike
Falls Church, VA 22043

Dr. Clinton Kelly
Decisions and Designs, Inc.
8400 Westpark Drive, Suite 600
McLean, VA 22101

Mr. George Pugh
General Research Corp.
7655 Old Springhouse Road
McLean, VA 22101

Mr. J. W. Stump
Grumman Aerospace Corporation
McLean, VA 22101

Mr. Gary W. Irving
Integrated Sciences Corporation
1532 Third Street
Santa Monica, CA 90401

Dr. Paul Slovic
Decision Research
1201 Oak Street
Eugene, Oregon 97401

Dr. Amos Freedy
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. A. C. Miller III
Stanford Research Institute
Decision Analysis Group
Menlo Park, CA 94025

Mr. Victor Rowney
Stanford Research Institute
Naval Warfare Research Center
Menlo Park, CA 94025

Dr. R. A. Howard
Stanford University
Stanford, CA 94305

Dr. H. L. Morgan
University of Pennsylvania
Wharton School
Philadelphia, PA 19174

Dr. Ward Edwards
Director, Social Science Research
Institute
University of Southern California
Los Angeles, CA 90007

Dr. Lee Roy Beach
Department of Psychology
University of Washington
Seattle, Washington 98195

Mr. Frank Moses
U.S. Army Research Institute
1300 Wilson Boulevard
Arlington, VA 22209

Robert G. Gough, Major, USAF
Associate Professor
Department of Economics, Geography and
Management
USAF Academy, Colorado 80840

Dr. T. Owen Jacobs
P.O. Box 3122
Ft. Leavenworth, Kansas 66027

Mr. R. Fratila
Naval Electronics Systems Command
Elex 330
Washington, D.C. 20360