Annual Technical Report PATR-1037-77-8
Contract F44620-76-0094
August 1977

ADAPTIVE ESTIMATION OF INFORMATION VALUES IN CONTINUOUS DECISION MAKING AND CONTROL OF REMOTELY PILOTED VEHICLES

RANDALL STEEB
KEN CHEN
AMOS FREEDY

Prepared For:

Air Force Office of Scientific Research
Bolling Air Force Base
Washington, D.C. 20333

PERCEPTRONICS
6271 VARIEL AVENUE • WOODLAND HILLS • CALIFORNIA 91367 • PHONE (213) 884-7470

NOTES

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

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

(18) (19) REPORT DOCUMENTATION PAGE	READ INSTRUCTIONS BEFORE COMPLETING FORM
REPORT NUMBER: AFOSR-TR-78-0193	2. GOVT ACCESSION NO. / 3. RECIPIENT'S CATALOG NUMBER
(6) 4. TITLE (and Subtitle): Adaptive Estimation of Information Values in Continuous Decision Making and Control of Remotely Piloted Vehicles.	(7) 5. TYPE OF REPORT & PERIOD COVERED: ~~INTERIM~~ Annual Technical rept.
	(14) 6. PERFORMING ORG. REPORT NUMBER: PATR-1037-77-8
(10) 7. AUTHOR(s): Randall Steeb, Ken Chen, Amos Freedy	(15) 8. CONTRACT OR GRANT NUMBER(s): F44620-76-C-0094 new
9. PERFORMING ORGANIZATION NAME AND ADDRESS: Perceptronics, Inc. 6271 Variel Avenue Woodland Hills, California 91367	(16) 10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 61102F 2313A4 (17) A1
11. CONTROLLING OFFICE NAME AND ADDRESS: Air Force Office of Scientific Research (NL) Bolling Air Force Base Washington, D.C. 20333	(11) 12. REPORT DATE: August 1977 / 13. NUMBER OF PAGES: 148
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) (12) 141p.	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)

Adaptive Models	Decision Making	Multi-Attribute Utility
Communications Models	Information Value	Remotely Piloted Vehicles
Computer-Aided Decisions	Information Seeking	
Decision Aid	Man-Machine Interaction	

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

This report describes research and development centered on evaluation of information needs in supervision of remotely piloted vehicles. The selection of information for transmission and display is a recurrent, subjective decision involving many factors - machine state, operator capabilities, communications costs, and channel limitations among others. An adaptive computer program has been developed which incorporates these factors into a multi-attribute decision model. The program is designed to capture the supervisory

390 285

operator's decision policy by using a training algorithm based on pattern recognition techniques.

Preliminary tests of the adaptive modeling approach were made using a task simulation resembling control of a remotely piloted vehicle. Individual subjects navigated the RPV through a changing, hazardous environment. In doing so, the operators selected different combinations of information and control allocation. The adaptive model was found to be more predictive of the subject's behavior than either a constant, unity weight model or an off-line method of weight estimation. Also, prediction of behavior increased with presentation of model-based recommendations to the subjects. Finally, the model was found to be useful in identifying differing decision strategies. The multi-attribute model thus formulated is expected to find application in evaluation of alternative information needs. Methods for management of communications by the remote element are also discussed.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

1. INTRODUCTION

1.1 Remote Systems Overview

Man and the functional environment with which he deals are becoming increasingly separate. Remote systems such as undersea robotic systems, remotely piloted aircraft, extraterrestrial teleoperator systems, and biomedical manipulators extend man's influence across distance, time, and function. But as these systems amplify man's capabilities, they simultaneously introduce new sources of uncertainty. The machine almost by definition acts as an intermediary between the task environment and the operator, making certain states inaccessible to direct observation. The operator must then depend upon artificial devices for sensing, processing, and communicating situational information.

To some extent, the amount of information exchange required is dictated by the degree of remote system autonomy. For example, operation of a simple remote effector without autonomous capability requies that the operator continuously close a real-time feedback loop around the remote element (Ferrell, 1973). Sufficient information must be transmitted to enable the operator to judge distances; interpret forms and shapes; appraise contacts, orientations, forces, and motions; and to issue complex commands (Bejczy, 1973). Much less information transmission is needed if the remote element takes over a portion of the routine, recurrent control and decision making functions. In this situation, the operator retains responsibility for evaluation, problem solving, and supervision, but is relieved of the continuous control function (Singleton, 1976). A good example of this form of supervisory interaction is advanced aircraft control. An aircraft control system has a hierarchy of control stages, some of which are delegated to the machine and some to the operator. The inner loops of aircraft control entail such functions as vehicle

stabilization and short-range guidance. These routine, high frequency tracking tasks are typically automated. The higher level loops of navigation and long-range guidance, involving the more discrete processes of problem solving and decision making, are normally delegated to the operator (Johansson, 1976). This allocation of function results in a decrease in the amount of continuous communications. Further, the allocation matches the information flow more closely to human information processing capabilities.

Remote systems demonstrate additional needs for minimizing communications, since the operator is physically separated from the system he is controlling. The information interchange in remote systems can be extremely costly and time consuming. Bejczy (1973) and Freedy, Hull, Lucaccini, and Lyman (1971) note that communications between an operator on Earth and a remote manipulator in space are limited by factors such as time delay, bandwidth, signal-to-noise ratio, and maximum video frame rate. Thus the remote operation becomes an increasingly laborious, slow, power consuming, and costly process as the distance to the remote element increases. Even in closely linked terrestrial applications, the communications costs can be high. Communications in remotely piloted vehicle (RPV) control demonstrate significant energy costs, reduced system reliability, and increased possibilities of detection (Fogel, Englund, Mout, and Hertz, 1974; Mills, Bachert, and Hume, 1975).

The substantial communication costs have led to development of greater autonomous capabilities in remote systems. Machine intelligence--denoted by Bejczy (1973) to be programmed control systems above the level of numerical (or memory) control--has probably received the most attention. Machine intelligence involves learning control based on pattern recognition, Bayesian estimation, reinforcement learning, stochastic approximation, and other sophisticated methods (see Fu, 1970, for a discussion of these techniques). The incorporation of machine intelligence allows the

operator to give more general commands. He can then rely on the remote unit to carry out the commands (Ferrell, 1973). Ideally, the necessary communications are thus reduced to status updates and interchanges of program objectives. In actuality, however, the operator must still frequently interrogate the system to verify performance and to identify critical situations.

One method of reducing such inefficient interrogation is to extend the function of the remote intelligence to include communications evaluation and management. This should be well within the realm of feasibility. Machine intelligence normally implies some explicit or implicit model of the remote element, its environment, and task objectives. The model should already be capable of assessing confidence levels for machine and human control in the immediate action. It should also be able to determine the efficacy of providing control information to the operator. In fact, some initial efforts have been made toward placing such a communications initiative with the machine. Information and control allocation techniques have been proposed using criteria based on queuing models (Rouse, 1975; Engstrom and Rouse, 1976), optimal control models (Sheridan, 1976), and multi-attribute decision models (Steeb and Freedy, 1976). The present work proposes to develop and integrate these efforts

1.2 Communications Evaluation as a Decision Problem

The communication of information between man and remote system is a good deal more complicated than deciding when to pick up the receiver. Choices must repeatedly be made regarding variables such as the mix of information sensing, processing, encoding, transmitting, and display. Throughout this process, a balance must be maintained between maximizing operator awareness of system operation and minimizing communications costs and operator load.

At first, this appears to be a mathematical optimization problem--like searching for a diet that provides the necessary nutrition at minimum cost. However, the mathematical programming techniques required for this optimization--linear programming, goal programming, dynamic programming, etc.--demand rigid problem structuring and continuous variables. More often, the communication decision is incompletely defined and involves choices among discrete rather than continuous alternatives. Thus the discrete operators used in decision theoretic techniques--matrices, difference operators, and detailed parameter enumerations--are more appropriate.

The specific decision to be modeled is what type of information to send and when to transmit it. The operator (or, equivalently, an information management system) must select an alternative on the basis of a number of multidimensional and risky (probabilistic) factors. Decision theory provides a normative framework for such choices, tying the decisions to situational factors (linear cue models) or to the impact of information on system effectiveness (multi-attribute expected utility models). Chapters 2 and 3 will provide a detailed comparison of these approaches.

1.3 Subjective Elements of Information Value

The form of the decision model can be developed from completely objective factors ("true" probabilities, dollar costs, etc.), from purely subjective factors (subjective probabilities, utilities for consequences), or from some combination of the two. A purely objective analysis of information value would have to include a complete mapping of system conditions and possible decision outcomes to operational goals. Felson (1975) states that only in a few highly structured situations can such an optimal model be derived. Also, even if such a model is developed, acceptance of the aiding provided to the operator may be

lessened since individual operator preferences would not be incorporated in the machine decisions (Hanes and Gebhard, 1966; Ferguson and Jones, 1969).

A more promising task is to elicit or infer the operator's goal structure and then to incorporate it in a model of the decision situation. This is the approach currently taken. By incorporating operator utilities in the model, complex evaluation and goal direction functions are performed by the operator, while normative aggregation functions are assumed by the computer. Loss of optimality in such a pragmatic approach should not be a major problem, as operator utilities for information have been found to approximate objectively derived values (Wendt, 1969; McKendry and Enderwick, 1971). Also, the subjective values may reflect aspects that are not analytically tractable at the present time. These aspects include timing factors and expected influences on subsequent information decisions (v.Winterfeldt, 1975).

1.4 Decision Modeling Philosophy and Objectives

One can use widely differing philosophies to model and aid the information-seeking decisions of the human operator. Available techniques focus on such diverse themes as uncertainty reduction, behavioral cue regression, and risky utility maximization. To a degree, all of these methodologies are potentially applicable to communications in remote systems.

Fortunately, there are some guidelines for model choice arising from the special circumstances of remote systems supervision. For example, behavior prescription by a normative model is far more important than simple prediction by a descriptive model, since the model is to be used for information evaluation and management. Also, the operator is

expected to explicitly consider the likelihood and importance of achieving specified system objectives. A normative model is again suited to capturing such goal-oriented behavior. Finally, incorporation of the model into the remote system demands simplicity, immediate access to model parameters, minimal interference with operator function, and generality. It is not guaranteed that any model will satisfy all of these criteria. Nevertheless, an attempt will be made to develop a normative model that satisfies many of the above goals.

The decision model has other potential benefits in addition to evaluation and management of communications. The model is expected to provide a framework for comparing alternative configurations of information sources, transmission systems, and displays. The model may also be used for selection and training of operators through comparison with "expert" judgment. Information needs may be disclosed through sensitivity analysis of model parameters. Finally, the model is expected to result in more consistent and effective operator decisions by providing on-line recommendations to the operator. Each of these possibilities will be developed in the following sections.

2. INFORMATION SEEKING BEHAVIOR

2.1 Overview

The activities surrounding the selection, acquisition, and processing of information can be surprisingly diverse. Even in the rematively specialized task of remote system supervision, information-related activities often constitute a majority of the operator's functions. The operator must maintain an awareness of the remote environment state, the machine state, the capacity and quality of the communication channels, and the progress toward objectives. This section will explore the more important techniques of modeling these inform tion-related activities. For simplicity of phrase in the analyses, the processes of information recognition, selection, and acquisition will all be subsumed under the term *information seeking.*

The major meth ologies of modeling information seeking behavior can be somewhat arbitrarily divided into three categories: informa ion theory, cue regression, and utility theory. Each of these techniques attempts to model the usefulness of information for the decision maker. The techniques differ in the amount of structure assumed by the decision model. Following a discussion of what constitutes relevant information (regardless of modeling philosophy), each of the methodologies will be discussed in turn. The chapter ends with a comparison of the potential contributions of the different techniques to remote systems analysis.

2.2 Information Relevance

Before bounding into the morass of mathematics, behavior, and models, it is necessary to define some conceptual rules. In particular, one needs to define what constitutes useful information and what properties

an information model should exhibit. These definitions can be derived from the basic relationship of information seeking to decision making.

In the most general sense, information has been described as data of value in decision making. The effect of the information is to reduce some element of uncertainty in the decision making process. The uncertainty may be concerned with the structure of the decision, or it may deal with the relations between the structural elements (Whittemore and Yovits, 1974; Nickerson and Feehrer, 1975). Specifically, Whittemore and Yovits define the structural elements of a decision to be the possible set of actions, outcomes, states of nature, and goals. Each of these may be defined along continuous scales or as discrete elements. The information may help to define the members or the domain of these sets. More commonly, however, the decision structure is already defined, and the information acts to define the relations between the structural elements. These relations are the parameters normally dealt with in decision analysis--the probabilities of the states of nature; the conditional probabilities of outcomes given certain actions; and the values of outcomes according to the goal structure. Information may contain structural data, relational data, or both. As a corollary, then, information is valueless in a non-probabilistic, completely structured decision.

It is not enough, however, that some uncertainty regarding the decision structure and relations is reduced, since behavior and hence, consequences, may be unchanged. To identify actually pragmatic or consequential information, Whittemore and Yovits (1974) define the informon--the minimum amount of information needed to change the state of a decision maker. The circumstances necessary to change this state have been best enumerated by Emery (1969): (1) the information must affect the existing representation of the decision situation; (2) the change in the representation must then affect the decisions made; and (3) an increase

must occur in the utility resulting from the changed decisions. Information thus has value only if it changes the organization's formal view of the world, if decisions are sensitive to such a change, and if the utility derived is sensitive to differences in decisions.

There are several reasons why information may not have this impact. Information may be ineffective in changing the situational representation and resulting decisions because the data is too coarse or too fine (Marschak, 1963). Information that is too coarse fails to distinguish between effectively different states of nature for at least one of the alternative actions. Information that is too fine differentiates between states having identical payoffs for all actions. Effective information--data that is not too fine or too coarse--is termed by Marschak to be payoff relevant.

In summary, then, a complete model of information seeking behavior must; (1) reflect the data's effect on uncertainty reduction, (2) identify the change in behavior, and (3) quantify the difference in decision consequences.

2.3 Information Seeking Models

2.3.1 General. The concept of evaluating information on the basis of its effectiveness in improving decision making has led to a variety of quantitative models for information seeking. The most influential are the normative models used for prescribing optimal behavior. These normative procedures involve maximizing gain or utility, minimizing losses, or achieving greatest uncertainty reduction. Unfortunately, the operator rarely acts optimally. More flexible and varied descriptive models are thus required to capture the information seeker's individual policy. Multiple regression, heuristic models, and modified normative models are

used widely for such descriptive modeling. The following review will attempt to characterize the various models in terms of normative qualities, descriptive capabilites, degree of completeness, and practicality.

2.3.2 Information Theory Models. Information theory provides a simplistic but direct means of specifying the impact or value of an item of information. The measure of information value is given in terms of reduction of uncertainty to a decision maker. In its early form, information theory (or more correctly, communications theory) used choice information as a measure of the freedom of choice one has in selecting a message from a population of symbols (Shannon and Weaver, 1949). In this limited application, the context, meaning, and effectiveness of the message are of no significance. The theory is concerned only with the probability of receiving any particular message for various conditions of the transmission system.

Whittemore and Yovits (1973) expanded the information theory methodology by redefining the concept of uncertainty and by introducing aspects of decision theory into the formulation. In their model, uncertainty may be associated with the structural aspects of decision making--the possible sets of actions, outcomes, state of nature, and goals--along with the relational connections between these structural aspects. Information is considered to reduce the uncertainty associated with either the structural or relational elements.

Just as the bit was developed as the primary unit in communications theory, a single index was formulated by Whittemore and Yovits (1973, 1974) to represent the impact of information. They derived an overall function of decision determinacy--the uncertainty surrounding the choice of a course of action. This measure represents the combined effects of the information

on the decision maker's understanding of the situation. Whittemore and Yovits represented this quantity in the following manner:

$$\text{Decision Determinancy} = \sum_{i=1}^{m} \left| p(a_i) - \frac{1}{m} \right| \qquad (2\text{-}1)$$

where

a_i, i=1, m are possible courses of action

$p(a_i)$ is the probability of choice of action i given the decision maker's knowledge

If each $p(a_i) = 1/m$, the situation is completely undetermined, and thus the above function is essentially a distance from indeterminacy.

It is assumed in this formulation that it is possible to obtain a distribution that reflects the operator's overall inclination toward the various courses of action. One possible form is the probabilistic model of expected utility. Here each action is selected with a probability related to its expected utility (Becker, DeGroot, and Marschak, 1963; also, see Section 2.3.3 for a description of expected utility models). In this way, the probabilities of choice are dependent on both the structural and relational components of uncertainty. The information results in changes to the model parameters, new courses of action become favored, and the decision determinacy is increased.

In general, this uncertainty reduction formulation accords with the requirments of responsiveness to changes in the situation and in the decision maker's objectives. It appears descriptive, but lacks a normative framework for directing behavior. Also, the formulation does not result in an easily derived scale of information value. The remaining techniques will be seen to be more definitive.

2.3.3 Utility Methods. Utility theorists provide an essentially optimal but not necessarily descriptive approach to modeling information seeking. They calculate the expected information impact on the estimate of the state of nature. Then the increase in decision effectiveness under the new situation estimate is determined. This provides a strong tie between information and action choices, but at the same time requires a well-structured decision task.

Development of a utility model of information seeking requires that the possible states of nature, information choices, actions, outcomes, and values can all be exhaustively enumerated. The values may be either objectively defined (i.e., costs, payoffs) or subjectively estimated (subjective values or utilities). Each action and state of nature is assumed to be associated with a payoff or utility, as shown in the payoff matrix in Figure 2-1 (Emery, 1969). Each utility value in this matrix represents a summation of the multidimensional consequences stemming from the specific action and state of nature.

A key component of the model is the state of nature during an action. This state is uncertain for two reasons (Emery, 1969; Sheridan, 1976).

(1) An action takes a finite time to implement, and so the states considered are future states. Such states are inherently uncertain.

(2) The states of nature are perceived only indirectly through an information system. Thus existing and future states are known imperfectly.

The information system can be characterized as an s by m Markov maxtrix relating messages to the possible unknown states of nature (Figure 2-2). The

STATE OF NATURE

ACTION		z_1	z_h	z_m
	a_1	$u(a_1, z_1)$	$u(a_1, z_h)$	$u(a_1, z_m)$
	a_k	$u(a_k, z_1)$	$u(a_k, z_h)$	$u(a_k, z_m)$
	a_x	$u(a_s, z_1)$	$u(a_s, z_h)$	$u(a_k, z_m)$

FIGURE 2-1. DECISION PAYOFF MATRIX

MESSAGE

STATE OF NATURE		y_1	y_j	y_n
	z_1	$P(y_1 \mid z_1)$	$P(y_j \mid z_1)$	$P(y_n \mid z_1)$
	z_h	$P(y_1 \mid z_h)$	$P(y_j \mid z_h)$	$P(y_n \mid z_h)$
	z_m	$P(y_1 \mid z_m)$	$P(y_j \mid z_m)$	$P(y_n \mid z_m)$

FIGURE 2-2. INFORMATION SYSTEM MATRIX

matrix elements $P(y_j|a_h)$ are the conditional probabilities of receiving message j if state h prevails. Off-diagonal elements indicate imperfections in the information system. Such an imperfect correspondence between state and message may be due to a failure of the messages to discriminate between states, or may be the result of random errors in the messages.

A Bayesian analysis can be used to determine the usefulness of acquiring information to sharpen the estimate of the state of nature (Wendt, 1969). Let $P(z_h)$ be the prior probability that state z_h will occur during the coming decision. The revised probability of state z_h, given the message y_j, can be given by Bayes' theorem:

$$P(z_h|y_j) = \frac{P(y_j|z_h) \cdot P(z_h)}{P(y_j)} \qquad (2\text{-}2)$$

where

$$P(y_j) = \sum_i P(y_j|z_h) \cdot P(z_h) \qquad (2\text{-}3)$$

The value of information derives from the fact that only one message, y_j, is selected from the set Y of possible messages, and that this message allows improved inferences about the state of nature. To determine the value of the information structure, then, one must estimate the probability and impact of receiving each potential message y. (The subscripts i, j, h will be dropped in the following discussion for simplification.) In the coming paragraphs, a decision policy will be established for action selection in response to the message received. Then an expectation taken over all possible messages provides a figure of merit for the information system.

The optimal decision rule for action selection in a non-conflict situation is the maximization of expected utility (EU). This rule specifies selection of the action with the highest probability-weighted utility. If a specific message y is observed, EU maximization taken the following form:

$$\alpha^*(y) = \max_a \sum_z P(z|y)\ u(a,z) \tag{2-4}$$

Thus the optimal decision rule $\alpha^*(y)$ selects the action a that maximizes the expected utility under the revised probability estimate $P(z|y)$. The value of an information system, $v(\Gamma)$, is calculated by summing across all possible messages:

$$v(\Gamma) = \sum_y \sum_z P(z,y)\ u(\alpha^*(y),z) \tag{2-5}$$

The joint probability of state z and message y can be decomposed in the following manner:

$$P(z,y) = P(z)\ P(y|z) \tag{2-6}$$

resulting in:

$$v(\Gamma) = \sum_y \sum_z P(z)\ P(y|z) \cdot u(\alpha^*(y),z) \tag{2-7}$$

The overall expected value of the actions taken using information system y can thus be calculated using three sets of parameters: (1) $P(z)$, the prior probability of each state; (2) $P(y|z)$, the information system matrix; and (3) $u(a,z)$, the utility or payoff matrix. The fair cost of an

inquiry with the information system is the difference in expected utilities with the without the inquiry:

$$\text{Fair Cost} = \sum_y \sum_z P(z)\ P(y|z)\ u(a^*(y,z)),$$
$$- \max_a \sum_z P(z)\ u(a,z) \qquad (2\text{-}8)$$

This analysis is suited for highly structured tasks. Not only must the possible states, messages, actions, and outcomes be specifiable, but the prior state probabilities and the conditional probabilities characterizing the information system must be derivable. The sequence of decision stages can be depicted using a decision tree, as shown in Figure 2-3. The tree is folded back by associating with each possible message the maximum expected utility of the subsequent actions. This folding back represents graphically the process of EU maximization. The favored information source I is then identified by comparing the expectations taken over all possible messages.

The state-message-action approach described above is directly applicable to simple, discrete information transfers. Examples of information sources exhibiting such a discrete nature are warning signals, status displays, and mode indicators. The more dynamic and multidimensional forms of information transfer such as video displays and radar scans, must first be partitioned into analyzable elements. The normally continuous states, messages, and actions must each be organized into a small number of meaningfully distinct catagories. This catagorization is typically based on payoff relevance. Catagories should be comprised of elements having equivalent consequences or implications.

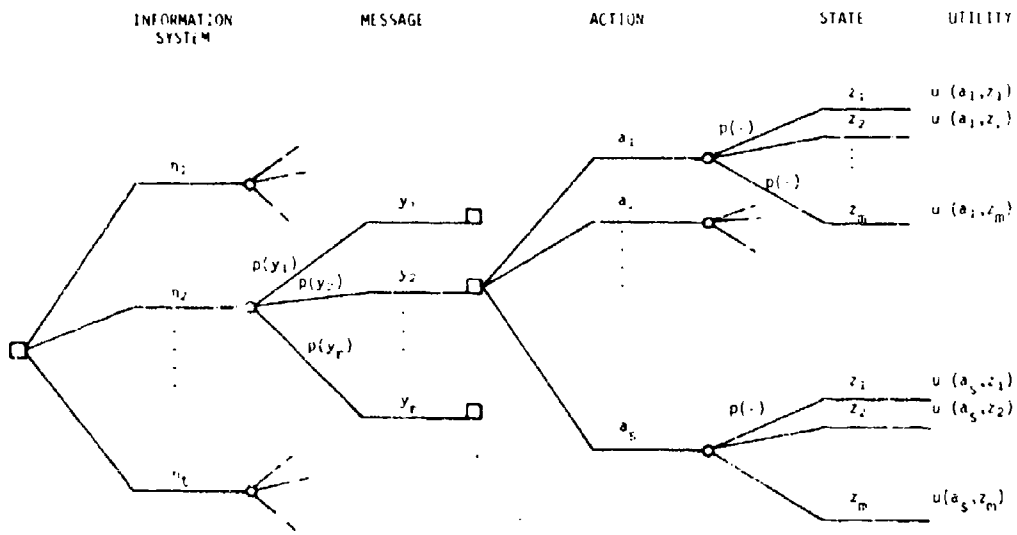

FIGURE 2-3. DECISION TREE FOR INFORMATION SYSTEM

For example, the environmental states and the resulting messages are often not difficult to discretize. The state variables by definition can be measured and described with only finite resolution. Also, the states are typically multidimensional, representing ensembles of environmental variables. Those elements of the states or messages that are irrelevant or unchanging can be ignored (Emery, 1969).

The possible actions can also be reduced to an elementary set. In addition to the formation of equivalence groupings according to consequence, there are techniques for identifying and deleting dominated actions. Actions which exhibit inferior consequences compared to some other action for every state of nature (dominated rows in Figure 2-1) can normally be deleted without problem.

Once the structure of the decision is established, the appropriate probabilities can be estimated. The key parameters are the conditional probabilities $P(y|z)$ characterizing each information source. These conditional probabilities can be estimated either objectively or subjectively. Objective estimation is simply the actuarial tabulation of the transmitted messages and the subsequently observed states. Subjective estimations are elicitations of judgments from the operator or inferences from his behavior. However, Slovic, Fischoff, and Lichtenstein (1977) state that man is a very poor Bayesian, systematically violating principles of rational decision making when dealing with probabilistic tasks. Similarly, Goldberg (1968), Rapoport and Wallsten (1972), and Beach (1975) have concluded that man's probabilistic judgments are unreliable across time and across diagnosticians, and only marginally related to his confidence in the accuracy of his judgments. In general, it appears that man is ill-suited for taking responsibility for complex probability aggregation. The methods of objective estimation are preferred whenever feasible. Once the conditional

probabilities $P(y|z)$ are defined either objectively or subjectively, the message probabilities [$P(y)$] and the updated state probabilities [$P(z)$] can be derived.

It should also be noted that an upper bound to the value of information can be easily calculated. This is the value of perfect information--the advantage provided by a completely clairvoyant information source. The perfect information tells with certainty what event will ensue. The difference in expected gain between an action taken with this information compared to an action taken without it provides a measure of the maximum possible value of information (Macrimmon and Taylor, 1972; Sheridan, 1976).

Utility Estimation. The utility $u(a,z)$ associated with a given state and action is typically a multidimensional quantity. The consequences associated with a given outcome may be characterized by such variables as resource expenditures, time delays, equipment losses, operator load, and goal attainment. For convenience, the consequence set is considered to include all costs and consequences, whether arising from the information acquisition or the action decision. This incorporation of information costs in the single function will simplify the parameter estimation processes in the coming sections.

The consequences may be either objectively or subjectively defined. For objective definition, the costs and gains associated with each outcome must be directly accessible. The choice of action can then be made solely on the basis of expected dollar return, ship-equivalents lost, or some other objective criteria. However, most real-life decision situations are so complex, unstructured, and poorly understood that such optimal decision systems cannot be designed (Felson, 1975). Instead, the operator's subjective value or utility for a given outcome must frequently be used as a guide. This is not necessarily bad. Operator utilities for outcomes or

information are often close to the objective values (Wendt, 1969; McKendry and Enderwick, 1971). Also, the subjective values may reflect important non-analytical aspects of outcomes, such as expected influences on subsequent decisions and individual operator needs. Finally, the subjective values are criteria that are condensed from experience, making the derivation of a complete but unwieldy dynamic model unnecessary.

The most tractable form for decomposition of the multidimensional outcomes is the additive multi-attribute model. This model requires satisfaction of several assumptions regarding behavior, which will be discussed in the context of remote systems decisions in Section 3.2. For now, let it be assumed that the axioms are met. The model takes the following form:

$$u(a,z) = u(x_1,x_2,\ldots,x_n) = \sum_{i=1}^{N} K_i u_i(x_i) \qquad (2\text{-}9)$$

where

x_i is the level of attribute x_i: $u(x)$ and $u_i(x_i)$ are utility functions scaled from 0 to 1

K_i is the scaling constant of attribute i

The notation used above and in the coming sections follows closely that used by Raiffa and Keeney (1975). Each distinct outcome is considered to be different combination of levels of the attributes x_i. Thus a single attribute vector characterizes each outcome.

The information seeking model (Equation 2-8) adds the aspect of risk to the multi-attribute formulation. Risk, in this context and in the remainder of this work, refers to a situation where the decision maker is able to specify a probability distribution over the possible outcomes of

an actio... Each choice in the decision tree of Figure 2-3 can, by folding back, be associated with an expected consequence vector. If axiomatically sound, this results in a multi-attribute expected utility formulation:

$$E\,[u(x)]_{\Gamma} = \sum_{J=1}^{M} P(x_j) \sum_{i=1}^{N} Ki\; u_i\,(x_{ij}) \qquad (2\text{-}10)$$

where

$E\,[u(x)]_{\Gamma}$ is the expected utility of inquiry Γ, and

$P(x_j)$ is the probability of outcome j with this inquiry

The normative nature of the EU model is well established, but its descriptive ability has been under a certain amount of attack. Tversky, Lichtenstein, and Slovik (1972) argue that descriptive models of choice must take into account cognitive variables such as memory and set. Similarly, Tversky and Kahneman (1973) have shown that decision makers often use heuristic, strain-reducing policies to simplify complex situations. In general though, the usefulness of EU models is conceded in situations where the number of attributes 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 behavior, unlike most of the other, heuristic-based models (Wendt, 1973; v.Winterfeldt and Fischer, 1973).

Less critical problems may also manifest themselves with the EU model. The analysis up to this point has been based on deterministic models of choice. The choice or action with the maximum expected gain is presumed selected, and randomness or change of behavior is not expressed by the model. A family of models that take such behavior randomness into account are the probabilistic models. These models are based on a theory of random preferences that can account for substantial errors or fluctuations

in behavior (v.Winterfeldt, 1975). Among the more important probabilistic models are the constant utility models (Luce and Suppes, 1965) and the random utility models (Becker, DeGroot, and Marschak, 1963). Constant utility models assume randomness in the response mechanism, while random utility models place this randomness in the utilities themselves. Unfortunately, the difficulties of assessing and using the probabilistic utility functions has made these techniques virtually intractable (v.Winterfeldt, 1975).

Similarly, utility models of dynamic decision situations (in which each decision affects the future decisions) are as yet unfeasible. Some dynamic programming models have used static expected utility measurements as inputs to their dynamic calculations (Slovik, Fischoff, and Lichtenstein, 1977). Thus far, however, no models have explicitly incorporated the dynamic nature of the decision environment into the utility measurements (v.Winterfeldt, 1975).

The techniques currently used for deterministic utility assessment can be divided into five main categories: ordinal scale methods, direct methods, gambling methods, regression techniques, and pattern recognition algorithms. The first three techniques have been thoroughly reviewed and analyzed by Kneppreth, Gustafson, Johnson, and Leifer (1973). With ordinal assessment methods, the decision maker is asked to qualitatively rank his preferences. His rankings are used to develop an ordinal scale of utilities. This can be converted to an interval scale if equal intervals are assumed, but the resulting sclae is only approximate.

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

One axiomatic procedure that is especially useful for quantifying utilities that vary on several value relevant attributes is conjoint measurement. This technique constructs a utility function over the multidimensional choice entities that decomposes to a single attribute utility function. One procedure prescribed for defining the functions involves indifference judgments to a sequence of unit steps in one attribute and varying amounts in a second attribute (v.Winterfeldt, 1975). While axiomatically valid, this process can be time-consuming and artificial to the decision maker.

Gambling methods require the *a priori* decomposition of complex decisions into many simple lotteries. Either the probability or the outcome of each lottery is varied until the decision maker is indifferent between the lottery and a "sure thing." Often, utilities associated with probabilistic outcomes are different from those found with riskless choices. Thus some form of gambling method is necessary for decisions under risk (Kneppreth, et al, 1973). Unfortunately, while the utilities thus calculated are axiomatically valid, the process is long, tedious, and somewhat contrived.

The regression and pattern recognition methods are the only techniques that estimate parameters from actual in-task behavior. These techniques assume a model of behavior, such as a multi-attribute EU model, and fit the parameters of the model to the observed behavior. The regression techniques usually require a large batch of observations and an interval scaled response. The model parameters are then estimated using a "least squared error" criterion. The pattern recognition approaches are more iterative than the regression methods. An initial set of parameter values is assumed and the model adjusts the parameter set decision-by-decision as incorrect predictions are made. The pattern recognition approach thus has the advantages of refining the model each time information becomes available, of requiring minimal memory, and of weighting recent observations more heavily (Pelson, 1975b; Weisbrod, Freedy, and Steeb,

in press). Also, the pattern recognition approach is very flexible in structure. Differing criteria of modeling performance--predictability, adaptability, robustness, etc.--can be met by varying the form of the model parameters and adjustment mechanism.

2.3.4 Cue Regression Approaches. Cue regression is a highly descriptive and pragmatic approach to modeling behavior. Rather than restricting the model to the limited realm of normative behavior, cue regression assumes only that the operator responds to situational cues in an algebraic fashion. Then the information choices are predicted by simple linear combinations of situational factors such as information cost, content, accuracy, and timeliness. These factors are the characteristics that contribute to the attractiveness of an information choice. The cues or features may also include the factors treated by the normative models--the specific outcomes of the decisions. However, the strict assumptions underlying the normative models are not considered. As a result, the cue regression approaches achieve great flexibility, but lack the goal-oriented power of the utility-based methods.

The parameters of the algebraic models are typically estimated through analysis of variance, conjoint measurement, and multiple regression (Slovik, Fischoff, and Lichtenstein, 1977). The simplest of these are the regression approaches. The regression methods employ correlational statistics to define a linear model of individual judgment. The linear model form was developed from Brunswik's (1940, 1952) famous lens model. This early model was an attempt to express the decision maker's policy of weighting various stimulus dimensions. Initially, Brunswik's cues and judgments dealt solely with information concerning the environment state, but later researchers expanded the model domain to include a wide range of judgments (Rapoport and Wallsten, 1972).

The structure of the cue regression model is quite simple. It is assumed that the decision maker provides numerical responses as judgments, and that the responses constitute some linear combination of the stimulus dimensions. In equation form, this becomes:

$$R = \sum_i W_i S_i + C \qquad (2\text{-}11)$$

where

R is the numerical response
S_i is the stimulus level on dimension i
W_i is the regression weight for dimension i
C is an optional scaling constant

The W's are regression weights reflecting the relative importance of each dimension. Estimation of the W's is accomplished by making the best fit of these weights to a batch of interval-scaled responses. The linear model thus developed is highly predictive if the predictor variables (the stimuli in Equation 2-11) are monotonic with the response function (Dawes and Corrigan, 1974). Unfortunately, the model may fail to suggest any underlying processes, as it is not axiomatic. In fact, the linear model is "paramorphic"--it does not presuppose the operator to additively consider the various stimulus dimensions. The model is simply predictive of the operator's choices.

The linear regression models have proved to be effective both in prediction of behavior and in replacement of the operator. Correlations between model-estimated parameter weights and subjectively elicited weights are quite high, normally in the .80 to .90 range (Dawes and Corrigan, 1974). Similarly, replacing the operator by a model derived from his previous judgments in the same situation (bootstrapping) is quite effective. Based on an autocorrelation model of the operator, these linear models often

perform better than the operators whom they model (Dawes and Corrigan, 1974; Slovik, Fischoff and Lichtenstein, 1977). That is, the correlation between the output of the model with a criterion (e.g., correct decisions) is often higher than the correlation between the operator's output and the criterion. The basis for this superiority of model over man appears to be the ability of models to eliminate or reduce "noise" effects or random behavior (Bowman, 1963). The applicability of these models is limited to situations involving recurrent decisions with relatively stationary behavior. However, such situations are common (Kunruther, 1969).

These policy-capturing approaches have been found by Dawes and Corrigan (1974) to be most effective in situations where (1) the predictor variables are monotonically related to the criterion (or can be easily rescaled to be monotonic), and (2) there is error in the independent and dependent variables. Dawes and Corrigan demonstrated that these conditions ensure good fits by the linear models, regardless of whether the weights in the models are optimal. In fact, Einhorn and Hogarth (1975) found that unit weights sometimes outstrip the estimated weights in predictive ability. They noted that unit-weighting schemes are effective in situations with errors in the model form, intercorrelations of variables, and small sample sizes. On the other hand, Newman (1975) states that unit weighting schemes are contraindicated in situations where there are negative correlations between attributes. He also notes that such circumstances are frequent. An applied comparison of the potential usefulness of unity and inferred weighting models will be found in Section 4.2.

2.4 Conclusions

The main considerations in the development of aiding in remote systems must be practicality and normative direction. Aiding in the information and control decisions faced by the human operator should be

based on optimal criteria of choice. Also, the aiding must rely on model parameters observable during task performance.

The information theory and cue regression approaches make only marginal contributions to this level of communications analysis. Neither approach provides a strictly normative basis for decision making. The entropy measures used in information theory do not identify the optimal information source or provide guidelines for choice. However, these methods do provide some direction when structural information is deficient. None of the other models appear to be able to deal with structurally incomplete decisions.

The linear cue models are well suited to prediction of information seeking decisions. These regression approaches can incorporate a variety of combinations of independent, monotonic situational cues to arrive at a predictive model. The attributes are not limited to decision consequences, and in fact, the ensuing action decisions do not even have to be considered in the information seeking model. This approach is preferred over the utility models if descriptive modeling of the information behavior alone is desired or if the action decisions are unobservable. Also, the linear cue models may be required if the axioms underlying the utility model cannot be satisfied.

The utility models provide the most powerful and normative approach. These models employ a multi-attribute, expected-utility formulation to model both the information seeking choices and the subsequent action decisions. The model is tied implicitly to system objectives since the model attributes are comprised of the constituent decision consequences. The information seeking is thus linked directly to the effect it has on augmenting system effectiveness.

With such a normative utility model, diverse aids are feasible. The communications configurations can be compared according to their contribution to the attainment of immediate system objectives. The operator's consistency between information seeking and action decisions can be ascertained and corrective feedback can be given. Also, the automated management of information can be based on the expected impact on action effectiveness, rather than on the simple mimicry of operator behavior provided by the cue regression approach.

Of course, the stronger implications of the utility model require more stringent assumptions of behavior. The axioms of both expected utility theory and multi-attribute aggregation must be satisfied. Also, the levels of the constituent decision consequences must be available as inputs to the model. As will be seen in the next section, these requirements are only rarely satisfied in a strict sense, but often can be realized to the necessary degree.

3. INFORMATION VALUE MODELING IN REMOTE SYSTEMS

3.1 Overview

The purpose of modeling is to provide a means of structuring, analyzing, and predicting the behavior of the system under study. The system in question, a human operator supervising a semi-autonomous remote element, is a very complex one. The objective of this work will be to examine a specific aspect of the remote system, the recurrent communications decision, and develop an elementary but tractable model. The key decisions of information seeking and control will be exhaustively structured, resulting in a definition of the set of actions, attributes, and consequences. A pattern recognition approach will be used to fit the model to observed behavior. In the end, the model should be capable of evaluating the information system and, to some extent, managing the communications according to operator needs.

Philosophically, the adaptive modeling pursued here will be closely related to the "on-line model matching" methods practiced in adaptive manual control (Baron, 1977). It will also be similar to the adaptive linear models used to augment or replace the expert decision maker (Bowman, 1963; Kunreuther, 1969; Dawes and Corrigan, 1974). These techniques assume the operator to respond consistently to situational circumstances and requirements. They then use pattern recognition, learning algorithms, or regression techniques to estimate behavioral parameters.

3.2 Structuring of Remote Systems Communications Decisions

The combination of supervisory human operator and remote system can be considered to be a partnership between two synergistic elements. Although a large amount of overlap in function occurs, each element contributes unique capabilities to the task. The machine generally

assumes responsibility for rapid, recurrent control functions, while the man tends to take over the supervisory functions--planning, problem solving, performance evaluation, etc. Of course, this mix is changing as higher level functions become attainable with automation.

The communications requirements between man and machine are tied closely to the functional allocation between the two. As the remote system becomes more competent and autonomous, the supervisory and control demands on the operator are lessened. This reduces the need for continuous communications, and in fact, changes the form of the directives transmitted. With greater machine capabilities, discrete statements of objectives and plan changes tend to be transmitted instead of continuous control commands (Johansson, 1976).

The basic form of the information flow between a remote system (here a remotely piloted aircraft) and a man is shown in Figure 3-1. The autopilot system is supplied continuously by onboard sensors with information regarding the machine state and environmental conditions. Assuming the remote element is equipped with a set of automated responses or learning algorithms, it should be able to respond autonomously to a variety of situational conditions. The human operator, on the other hand, must rely completely on the communications channel for information and control. His only access to the environmental conditions and machine state is through the communications interface. In the same way, the operator must rely on the communication interface as a means of transmitting commands and objectives to the remote system.

A learning control system is ideally suited to the modeling of information seeking decisions. Almost be definition, a learning system must partition into meaningful catagories the possible states, messages, and actions. Also, learning systems often employ conditional probability models for determining action policies (Freedy, 1969; Kanal, 1974). Thus

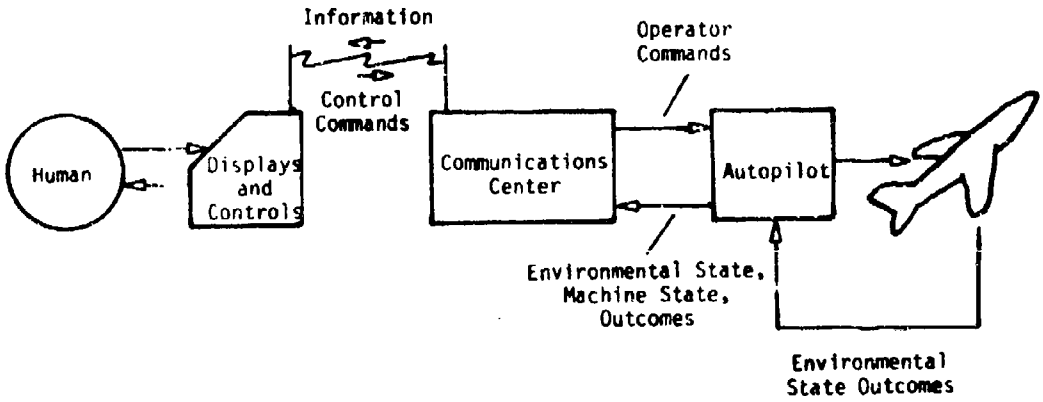

FIGURE 3-1. STRUCTURE OF MAN-MACHINE INTERACTION

the same prior probability estimates and conditional probabilities described in Section 2.3.3 for the information seeking model may already be resident in the autopilot program. Transfer of the information seeking model to such a learning system should be straightforward.

Communications between the remote system and the human operator, on the other hand, are much more difficult to model using a state-message-action paradigm. The information transmissions provided to the operator--video displays, radar sightings, infrared scans, etc.--are often complex, dynamic, and multidimensional. The operator normally does not respond to these information displays as discrete messages, but rather as dynamic, pictorial displays. This situation can be illuminated by looking at the observable inputs and outputs of the human operator and of the remote autopilot. These inputs and outputs are diagrammed in Figure 3-2. The autopilot, being contiguous with the remote environment, has access to the sensed conditions, the consequences sustained, and various operator inputs. The autopilot uses these inputs to estimate the information system characteristics, calculate the expected consequence levels for combinations of information and control, and select optimal actions. Also, as will be shown later, the remote system will often be able to use the operator inputs to infer the operator's value structure. The human operator, shown in the lower portion of Figure 3-2, has a different set of inputs and different responsibilities. He is appraised of the communications choices open to him and of their expected consequences. He may request an information transmission or he may be provided one automatically by the remote element. His observable outputs are the information requests and control commands.

The problem, then, is to develop an information and control model which takes into account both the discrete, deterministic nature of the remote automaton and the complex, intuitive processes of the human. An

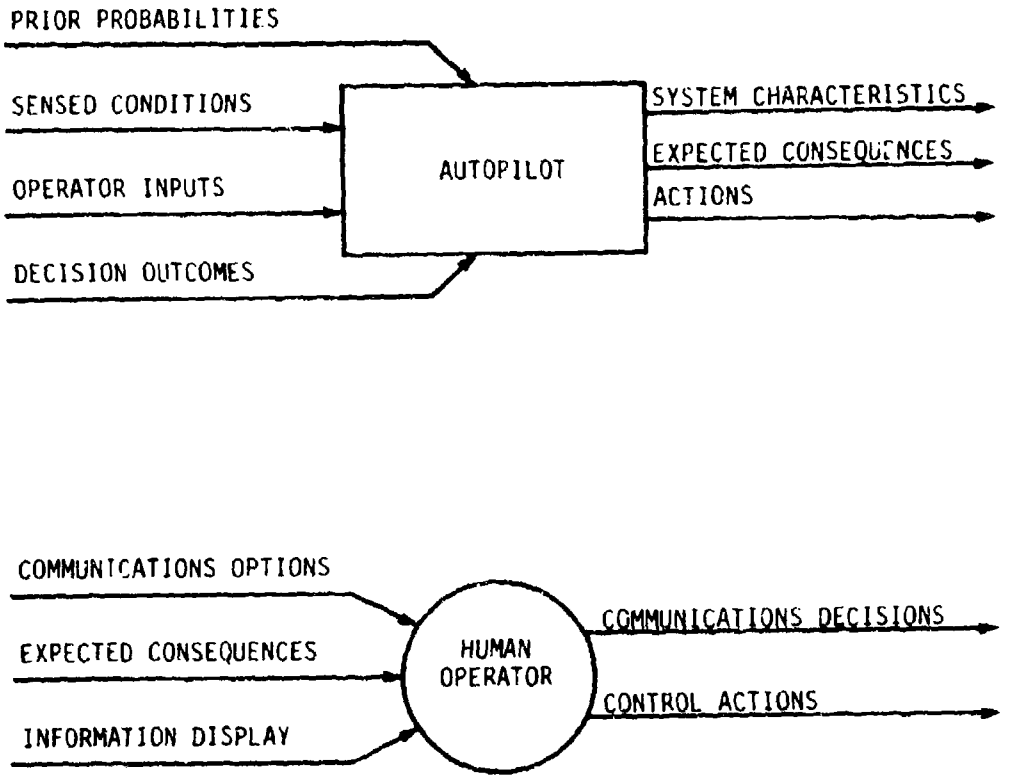

FIGURE 3-2. OBSERVABLE SYSTEM INPUTS AND OUTPUTS

attempt to diagram such a decision structure is shown in Figure 3-3. Two stages are evident, just as in the information seeking model depicted previously in Figure 2-3. The first stage concerns the sensing of the environmental state by the remote element. Information sources are selected according to their costs and their impact on the prior probability estimate $P(z)$. A given information source results in a sampling of a set of possible messages. The probability of a given message y depends on the prior probability $P(z)$ and on the information system characteristics $P(y|z)$. Once a message is received, the autopilot updates its state estimate and selects an appropriate action.

The second stage, action selection, is the critical one with respect to the man-machine interaction. The autopilot must respond to the apparent circumstances by either selecting a direct control action or by opening the communications channel to the operator. The two actions--autopilot control versus delegation of information and control to the human operator--are treated quite differently. The autopilot control action is considered to be an optimal response to an uncertain state estimate. The outcome is a deterministic function of the action and the true state z. The costs incurred during acquisition of the information are also included in the consequence set. The allocation of information and control to the operator, on the other hand, is treated as an information transmission with only partially observable parameters. The channel opening is simply considered to be an action with a distribution of outcomes which depend on the true states. Section 3.3.2 will develop these concepts more carefully.

The outcomes or vectors of consequences seen at the right of Figure 3-3 are all of the same space of dimensions, whether for machine actions or operator control. The levels of the consequences will be either

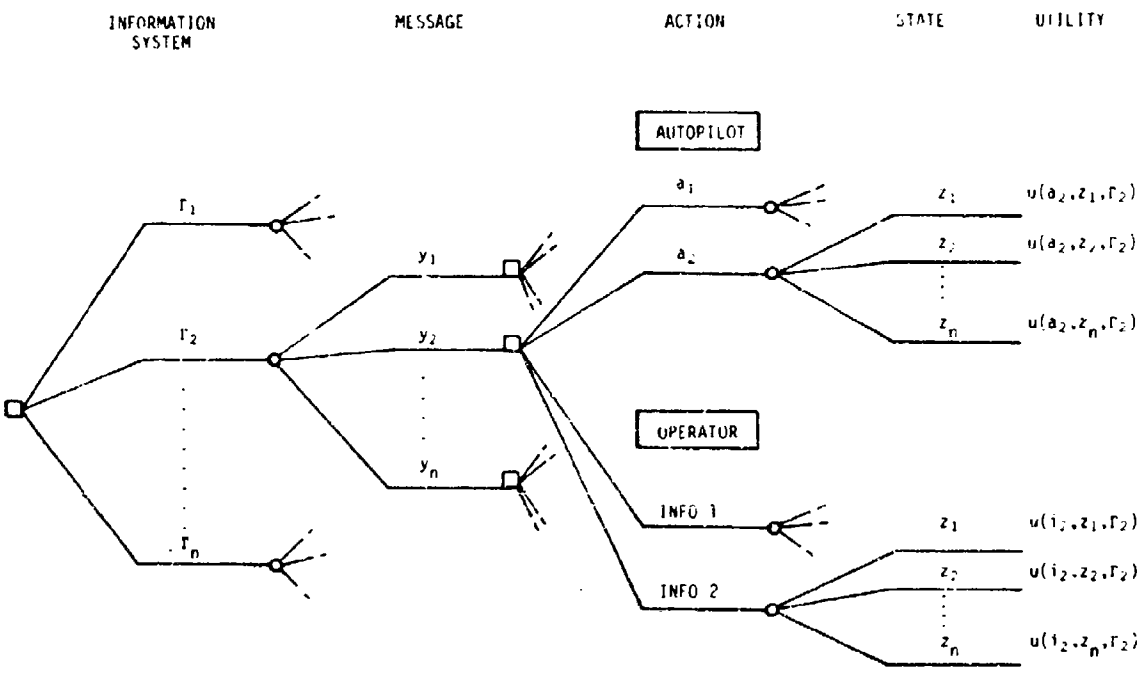

FIGURE 3-3. DECISION TREE FOR INFORMATION SEEKING DECISION

input prior to system operation, or estimated using predictive features and performance histories. The utilities for the consequence vectors will be elicited from the operator or inferred from his behavior.

The decision diagrammed in Figure 3-3 is by necessity a simplification of reality. For example, the basic two-stage structure does not consider the continued sampling of information prior to an action or the possible generation of new alternatives due to the information. Including such factors into the model at this stage would make it excessively complex. The remainder of this chapter will analyze methods of modeling the two-stage communication decision. The modeling will include both the structuring of the decision and the means of estimating model parameters.

3.3 Multi-Attribute Decision Model

3.3.1 General. Several steps are necessary to define the information and control model outlined in Section 3.2. First, the various event probabilities $P(y)$, $P(z)$, and $P(z|y)$ must be estimated or ascertained by observation and adjustment. Then the consequences associated with each outcome must be scaled along a set of dimensions. The decision tree is then folded back, associating an expected consequence vector with each information and action choice. Finally, the subjective importance weights associated with each consequence dimension are determined by elicitation or by inference from behavior.

These processes are based on the three major normative theories of decision making--Bayesian revision of probabilities, expected utility maximization, and multi-attribute utility analysis (Pitz, 1975). For these methods to apply, a variety of assumptions concerning the decision behavior must be satisfied. This section will explore the implications of these

assumptions. The chapter will also develop procedures for specification of the decision parameter set and estimation of the parameter levels.

Before continuing, some clarification of notation needs to be made. The system to be used will, for the most part, follow that of Keeney and Raiffa (1975). The following terms from the core of this notation:

Attributes: The basic attributes will be $X_1, X_2, \ldots, X_n$, where X_i may be either vector or scalar.

Attribute Sets: A complete set of attributes is defined as $X = \{X_1, X_2, \ldots, X_n\}$. A subset Y of X may be defined by identifying the attributes X_i in the subset.

Consequences: The consequence space $X_1 \times X_2 \times \ldots \times X_n$ represents a Euclidean Space. Consequences are designated by $x = (x_1, x_2, \ldots, x_n)$ where x_i corresponds to a specific amount of X_i for $i = 1, 2, \ldots, n$.

Relations: Preferences among consequences are denoted by the following relations: $\geq$ indicates preferred over equivalent to; $\sim$ indicates equivalence.

Scaling: The symbol $x^* = (x_1^*, x_2^*, \ldots, x_n^*)$ represents the most desirable consequence, and $x^\circ = (x_1^\circ, x_2^\circ, \ldots, x_n^\circ)$ designates the least desirable. The utility function of x (with the appropriate assumptions satisfied) is scaled by assigning $U(x^\circ) = 0$ and $U(x^*) = 1$.

Probabilities: The probability of occurrance of an event z will be denoted by P(z); a joint probability of a and z will be P(a,z); a conditional probability of y given z will be P(y|z).

Risky Outcomes: An option of receiving x' with probability Π_i and x" with probability $1-\Pi_i$ is expressed by <x', Π_i, x">. The superscript prime denotes a distinct variable, not a derivative. A risky (probabilistic) variable is designated by a tilde: $\tilde{x}$.

3.3.2 Probability Aggregation. The probability calculation procedures are centered around Bayes' rule expressed earlier in Equation 2-2. The use of Bayes' rule presupposes several key assumptions. Most importantly, the hypothesized states of the world must be exhaustive and mutually exclusive (Nickerson and Feehrer, 1975). Each state z_h is assigned an *a priori* probability of occurrence $P(z_h)$. Because these probabilities are mutually exclusive and exhaustive, they sum to one:

$$\sum_h P(z_h) = 1 \qquad (3\text{-}1)$$

Data or messages concerning the state may be used to revise the *a priori* distribution if the data is in the form of discrete, conditionally independent observations. Conditionally independent messages are those which are dependent of each other with respect to the states. This consideration of the states make conditional independence a stronger assumption than simple independence. Beach (1975) gives the following example showing the difference between the two conditions:

> *"...suppose medical research shows that across diseases there is no relationship between blood pressure and fever. These two symptoms would then be considered independent. It is possible, however, that there is a disease or a set of diseases for which blood pressure*

and fever are related, i.e., that they are conditionally dependent. Use of Bayes' theorem with these diseases as hypotheses is inappropriate because the redundant information or the extra information latent in the combination of the two symptoms (comparable to an interaction effect) can result in higher or lower posterior probabilities than is appropriate."(p.44)

The notion of conditional independence appears to be of importance only when messages are selected in groups. When making revisions based on single messages, the set of messages from which the datum is sampled need not be independent. Even if multiple messages are present, Beach (1975) notes that methods are available for circumventing the requirements. Conditionally dependent data can be dealt with using a modified form of Bayes' theorem, or the data can be "chunked" into conditionally independent groups.

The sensor characteristics $P(y|z)$ can be derived from observation. Comparisons of the messages received and the states subsequently observed provide the necessary data. Estimates of $P(z|y)$ and $P(y)$ result from frequency counts. $P(y|z)$ then can be calculated using the following expression:

$$P(y|z) = P(z|y)\ P(y) \tag{3-2}$$

In combination with the prior probability estimates, the sensor characteristics allow for the estimation of $P(y)$, the probability distribution of each specific message.

$$P(y) = P(y|z)\ P(z) \tag{3-3}$$

Finally, when a particular message y is received, an updated P(z) can be calculated:

$$P(z|y) = \frac{P(y|z)\ P(z)}{P(y)} \qquad (3\text{-}4)$$

The above calculations are applicable to the decisions surrounding autopilot control. As mentioned earlier, the analysis of the human control branch of Figure 3-3 is somewhat less deterministic. It is true that with human control, the states, information system characteristics, and consequences are observable entities, just as with autopilot control. However, the information provided the operator is seldom amenable to decomposition into discrete messages. The rich mediums of video and radar display are typically too complex and multidimensional to be decomposed into such an analytical formulation. Consequently, the updated state probability estimate contingent on the message will also be unobservable. This state estimate is within the mind of the man. In fact, if a complete message-state-action specification were possible, then the entire cycle should probably be automated.

Nevertheless, the provision of a specific information transmission to the operator is an action with an observable set of consequences. The consequences depend on the true state of the world, just as they do for autopilot control. The action/state combination differs from that under autopilot control because the outcome is not deterministic--a variety of outcomes may occur depending on the specific action taken by the operator. Thus an expected consequence vector, estimated from a number of observations, is associated with each action (form of information transmitted to the operator) and state.

The above procedure essentially results in a collapsing of one stage of the decision tree. Figure 3-4 details the processes resulting

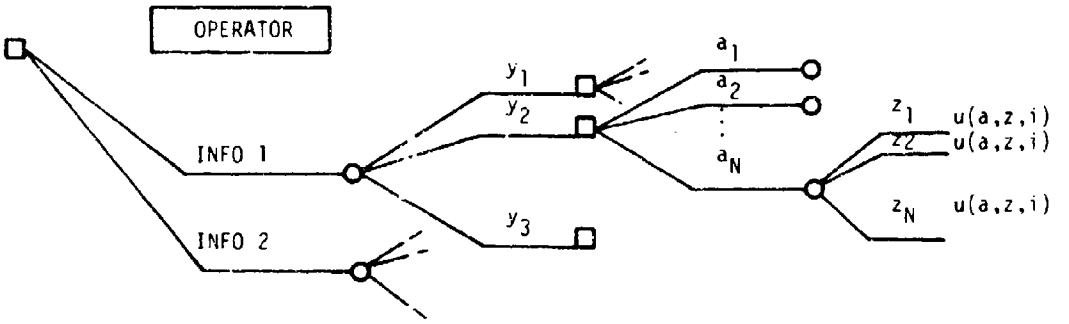

a) COMPLETE SPECIFICATION

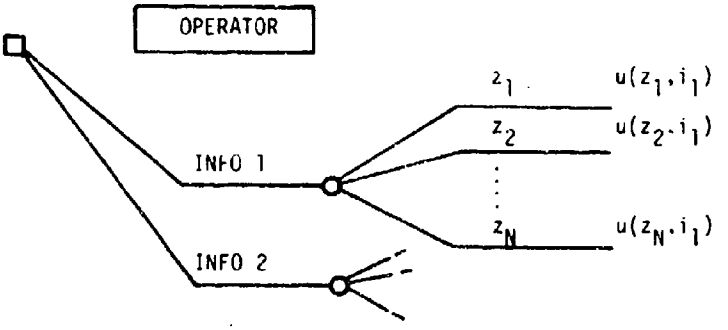

b) OBSERVABLE LEVEL OF SPECIFICATION

FIGURE 3-4. COMMUNICATIONS DECISION

from opening the communications channel to the operator. This is an elaboration of the lower right portion of Figure 3-3. The operator receives a specific but complex message from one of the information sources, updates his estimate of the situation, selects an appropriate action, and obtains an outcome dependent on the true state of the world. The lower figure, representing the observable stages, compresses the sequence to one of action and outcome. The outcomes for a given state in the lower figure are simply the expectation of all the outcomes for the corresponding state shown in the upper figure.

The probability estimates $P(z)$ for human control are the same as those for machine control. The updated state probability $P(z|y)$ is that derived from the initial autopilot receipt of sensor data. As mentioned earlier, the simplest method of determining the conditional probabilities is to maintain frequency counts of the various messages and subsequently observed states. If the conditional probabilities vary in time, then moving averages employing an observation window of a set number of past decisions can be used. Also, an exponential weighting of past decisions can be used to provide a bias adding to the importance of recent observations.

If these objective methods cannot be used, subjective probabilities can be elicited from the decision maker. A variety of approaches are available for such expressions (see for example, Goodman, 1973). Such elicitations add an additional subjective element to the model, making it the more descriptive, subjective expected utility (SEU) form. Of course, such subjective elicitations again have the drawback of interfering with the operator's task. Objective estimates are preferred because they tend to be more accurate and less burdensome.

3.3.3 Utility Analysis. The set of consequences associated with each action and state must be evaluated along well-defined interval scales. The relative importance of resource expenditures, time delays, vehicle losses, operation attention, and other consequences of operation can be either objectively defined or subjectively determined. Objective definition entails development of a mapping of organization goals to specific consequences. Such a mapping is extremely difficult to realize in complex situations. Such factors as future consequences, interactions, and subjective needs are often virtually undefinable. In situations involving such factors, it appears more useful to elicit or infer from the operator his utilities for the consequences.

The technique of greatest potential for providing a framework for subjective evaluation is that of multi-attribute expected utility theory. This is a relatively new methodology designed explicitly for complex decisions under risk. The analysis that follows is primarily attributable to Raiffa and his colleagues (Raiffa, 1969; Keeney and Raiffa, 1975), and to V.Winterfeldt and his associates (v.Winterfeldt and Fischer, 1973; v.Winterfeld, 1975). The intent of this analysis is to provide an axiomatic basis for the additive expected utility representation:

$$U(\tilde{x}) = \sum_{h=1}^{m} P(z_h) \sum_{i=1}^{n} U_i\ (x_{ihk}) \qquad (3\text{-}5)$$

where

$\tilde{x}$ represents the vector of consequences of risky action k

$P(z_h)$ is the probability of state h

x_{ihk} is the level of attribute i associated with state h and action k

U_i is the utility function over the ith attribute

The expression in Equation 3-1 is a decomposition of the decision into actions, event states, and attribute levels. The level of decomposition possible depends on several crucial independence assumptions. These assumptions derive from the origin of the model as a combination of expected utility maximization and multi-attribute utility aggregation. Thus the respective axiomatic treatments of these two methodologies must be satisfied. For expected utility, v.Winterfeldt and Fischer (1973) note that the theories of von Neumann and Morgenstern (1947), Savage (1954), and Luce and Raiffa (1957) all make two central assumptions concerning preferences among risky choices:

(1) Sure thing principle. Preferences among risky alternatives should be independent of events in which these alternatives have common outcomes. For two events, this is expressed by:

$$\langle x, \Pi_1, y\rangle \geq \langle z, \Pi_1, y\rangle$$

if and only if

$$\langle x, \Pi_1, w\rangle \geq \langle z, \Pi_1, w\rangle$$

(2) Solvability. No outcome should be infinitely desirable or undesirable. Thus, for all outcomes x, y, and z for which $x \geq y \geq z$, there is a Π_1 such that

$$y \sim \langle x, \Pi_1, z\rangle$$

These axioms allow formation of the single attribute expected utility function:

$$U(x) = \sum_{h=1}^{m} P(z_h)\, U(x_{hk}) \qquad (3\text{-}6)$$

Multi-attribute EU models can be developed from these axioms in two ways, transformation and decomposition. The less rigorous, transformation method uses riskless multi-attribute assumptions to construct a riskless utility function. This function is then transformed into a risky function using the expected utility formulation above.

The riskless MAUT model requires that preferences for values of each attribute are independent of constant values in the other attributes. Consider a set of attributes X partitioned into an arbitrary subset Y and its complement $\overline{Y}$. Then for any riskless consequences y', y", $\overline{y}$, $\overline{y^+}$

$$u(y', \overline{y^+}) \geq u(y'', \overline{y^+}) \Rightarrow u(y', \overline{y}) \geq u(y'', \overline{y}) \qquad (3\text{-}7)$$

This is termed weakly conditional utility independence (WCUI) by Raiffa (1969) and by v.Winterfeldt and Fischer (1973), preferential independence (Keeney and Raiffa, 1975), and single calcellation. If the test is satisfied for all attributes, then a riskless combination model is justified. The riskless model is generalized to include uncertainty by defining an expected utility model with a function U defined over multi-attributed alternatives. This is possible only if U is a linear function of the attributes (v.Winterfeldt and Fischer, 1973). A possible representation of this function is:

$$U(\tilde{x}) = \sum_{h=1}^{m} P(z_h)\, u(x_{hk}) \qquad (3\text{-}8)$$

in which

$$U(x_{hk}) = \sum_{i=1}^{n} W_i \, U_i \, (x_{ihk})$$

where W_i is a scaling constant for attribute i

A second, more restrictive approach to modeling risky, multi-attributed choices is the decomposition technique. This method first constructs the utility function U and then adds assumptions justifying the decomposition of U into individual components (v.Winterfeldt and Fischer, 1973). The main test is that of strong conditional utility independence (SCUI). This axiom states that preferences among risky alternatives, in which a subset of the attributes has constant values across the outcomes, should not depend on these constant values. Explicitly, for any lotteries $\tilde{y}'$ and $\tilde{y}''$, for any riskless consequence $\overline{y^+}$, and for all $\overline{y}$:

$$u(\tilde{y}, \overline{y^+}) \geq u(\tilde{y}'', \overline{y^+}) \Rightarrow u(\tilde{y}', \overline{y}) \geq u(\tilde{y}'', \overline{y}) \qquad (3\text{-}9)$$

Satisfaction of this assumption implies that the model form is either additive or multiplicative. A second, stricter axiom guaranteeing activity is that of marginality (Raiffa, 1969) or additive independence (Keeney and Raiffa, 1975). Marginality requires that the alternatives are judged solely on the basis of the marginal probability distribution over single attribute values. This implies:

$$\langle (y', z'), (y^\circ, z^\circ) \rangle \sim \langle (y', z^\circ), (y^\circ, z') \rangle \text{ for all } y', z' \qquad (3\text{-}10)$$

The marginality condition seldom holds in practice because of operator preferences concerning the variance of outcomes (see Wendt (1973)

for a discussion of variance preferences). If SCUI and marginality do hold, an additive expected utility model of the following form is justified:

$$U(\tilde{x}) = \sum_{h=1}^{m} P(z_h) \sum_{i=1}^{n} U_i (x_{ihk}) \qquad (3\text{-}11)$$

This is a risky decomposition model, compared to the previous riskless transformation model (Equation 3-6). The difference lies in the presence of weighting constants in the transformation model. This difference is not important in the present context, since both additive models will be seen to be amenable to the adaptive techniques of estimation developed in Section 3.4.

Strict adherence to the axiomatic assumptions is apparently not crucial. Experience has shown that models and procedures with differing axiomatic backing will produce convergent utility functions in a large number of cases (v.Winterfeldt, 1975). For instance, in *riskless* decision making, Yntema and Torgerson (1961) and Fischer (1972) demonstrate that additive models can approximate non-additive models quite well. Riskless linear models in regression also produce good results when compared with more complex models that include interactions (v.Winterfeldt and Fischer, 1973). Finally, Fischer (1972) showed that variations in the shape of riskless single attribute utility functions will produce overall utilities that are highly correlated as long as all single attribute functions are monotonic.

Specific studies of the importance of axiomatic backing in *risky* multi-attribute models are more difficult to locate, apparently because of the complexity of assessment required. v.Winterfeldt and Fischer (1973) recount that in v.Winterfeldt's (1971) dissertation, the risky MAUT measurement procedures were not adversely affected by axiomatic lapses.

Direct tests of the independence assumptions (Equations 3-7, 3-8) showed satisfaction of the SCUI condition but violation of marginality. Nevertheless, a correlational analysis indicated that the decomposition model was still quite effective. Fischer (1972) made similar observations studying preferences for risky job alternatives described by three attributes.

The observed insensitivity to axiomatic strictness may be related to the complexity of the decisions. When the number of dimensions is not large, v.Winterfeldt and Fischer (1973) say 5 or less, operators are fairly consistent in behavior and the models closely reproduce the operators' holistic responses. This is evocative of Miller's (1956) findings that people can deal with only 5 or 10 "chunks" of information at a time. Consequently, the random error (and model insensitivity) associated with a decision tends to increase as the DM attempts to consider an increasing number of value attributes (Slovic and Lichtenstein, 1971).

The axiomatic procedures are also subject to problems of assessment. Often, the tests are impossible to complete in complex real choice situations. None of the axioms can be verified absolutely, since they normally apply to an infinite domain. Also, the axioms require judgments that the decision maker is generally unable to make, such as ordering complex alternatives consistently (v.Winterfeldt and Fischer, 1973). Consequently, the applied work has been more concerned with structuring decision problems, assessing model parameters, and making sensitivity analyses, than with axiomatic tests. The axioms simply provide a tool to eliminate models that are clearly wrong. The basic axioms can be tested roughly by presenting the decision maker with "easy" choices. In this way, the DM should violate model assumptions systematically, so that it can be discovered which assumptions are appropriate.

3.3.4 Model Form. The multi-attribute EU model assumes that the decision outcomes are an amalgamation of many factors, each contributing to the overall attractiveness of the outcomes. The most widely used means of combination of the attributes are the additive, the multiplicative, and the multilinear forms (Keeney and Raiffa, 1975; Fischer, 1972). The respective expressions are:

$$U(x) = \sum_{i=1}^{n} K_i U_k(x_i) \qquad \text{(Additive)} \qquad (3\text{-}12)$$

$$1+KU(x) = \prod_{i=1}^{n} [1 + KK_i U_i(x_i)] \qquad \text{(Multiplicative)} \qquad (3\text{-}13)$$

$$U(x) = \sum_{i=1}^{n} K_i U_i(x_i) + \sum_{j>1} K_{ij} U_i(x_i)\, U_j(x_j) + \ldots + K_{1\ldots n}\, U_1(x_1) \ldots U_n(x_n) \qquad \text{(Multilinear)} \qquad (3\text{-}14)$$

where the K_i, K_{ij}, ... $K_{1\ldots n}$ are scaling constants, $0 < K_i < 1$ and $K > -1$ is a non-zero scaling constant satisfying

$$1 + K = \prod_{i=1}^{n} (1 + KK_i)$$

The multiplicative and multilinear forms are needed if the factors are not treated in an entirely compensatory fasion. A compensatory model is one in which changes in one attribute can compensate for changes in another. The multiplicative and multilinear forms can be non-compensatory since an extreme value of one factor allows it to dominate all other factors. Also, these non-compensatory models can account for some configural effects--interactions and higher order terms.

The additive form, while simply a special case of the multiplicative form (Keeney and Sicherman, 1975), appears to be the one best suited for modeling and aiding. The linear form of the additive function is more appropriate for estimation by pattern recognition techniques than the complex structures of the multiplicative and multilinear models. In fact, Section 3.4 will demonstrate how the linear additive form can be used directly as a discriminant function. The additive model is also more suitable for use in analysis and feedback of behavioral characteristics. For example, the use of outcome probabilities as model parameters makes sense only for the additive model (Huber, 1974). Huber also notes that the additive model is more robust to unsatisfactory attribute levels than the multiplicative model. An erroneous zero in one of the additive factors does not have the major effect seen in a multiplicative model.

In certain cases, the additive MAU model may even be modified to account for interdependencies among factors. Keeney and Sicherman (1975) define a special "nested" attribute that consists of a vector of sub-attributes. This factor provides an extra degree of freedom through an extra scaling constant. Thus trade-offs between two factors can depend on a third. This reduces the need for the configural terms provided by the multiplicative and exponential models.

The favored model form is thus a weighted sum (Equation 3-8) of outcome components, some of which may be probabilistic. This model is postulated to be applicable to both stages of the remote systems communication decision--the information seeking decisions and the ensuing action selection choices. Both of these decisions entail evaluations of the same set of attributes. Also, the same weighting factors w_i are expected to apply equally to the two types of decisions.

3.3.5 Factor Development. The choice of factors to include in the decision model is probably of greater importance than the choice of the model form itself. This is evident from the frequent effectiveness of unity (arbitrary) weighting schemes in predicting choices. Dawes (1975) states: "The whole trick is to decide what variables to look at and then know how to add." Unfortunately, guidelines for the choice of model attributes are not readily available. The following list of desirable characteristics expands on Raiffa's (1969) recommendations of attribute independence, set completeness, and minimum dimensionality:

(1) Accessible. The levels of each factor should be easily and accurately measurable.

(2) Conditionally Monotonic. The factor level should be monotonic with the criterion (preference) regardless of the constant values of other factors.

(3) Value Independent. The level of one attribute should not depend on the levels of the other attributes. This is to some extent a consequence of recommendation number two.

(4) Complete. The set of attributes should account for as much as possible of the operator's behavior.

(5) Meaningful. The attributes should be reliable and should demonstrate construct validity. Their implications should be understandable when expressed in feedback to the operator.

For the most part, these recommendations result in an attribute set that is accessible, predictive, and in accord with the axioms of utility theory. The recommendations also imply a limitation on the number

of possible attributes. The requirements of independence and meaningfulness render any large set of attributes unrealizable, because of the cognitive limitations of the human operator.

A candidate set of attributes x_i for the remote systems decision task could include such factors as resource expenditures, time usage, vehicle losses, operator supervisory load, and future consequences. The levels of each of these attributes can be derived using relations between the attributes and such situational features as environmental conditions, communications channel characteristics, and autopilot capability. The relations may entail probabilistic mappings from features to attributes, or they may involve simple transformations. In the absence of available situational features and mappings, subjective estimates can be used (Edwards and Gutentag, 1975). Of course, the expression of such subjective estimates may be burdensome or may require costly communications between the operator and the remote system. Thus attributes which are measurable by the remote element are favored.

An attribute requiring special attention is cost. Information acquisition costs are often considered separately in decision models. However, a variety of types of costs may be incurred--energy costs, equipment usage, risk of detection, etc. Thus appears more logical to associate a vector of incurred costs with the respective outcome. Also, if cost is subject to an upper bound due to limited resources, the ratio $U(x)|c$ for each alternative must be calculated (Edwards and Gutentag, 1975). These are the famous benefit to cost ratios. Actions should be chosen in decreasing order of that ratio until the budget constraint is used up. In the absence of a budget constraint, cost is just another additive dimension of value.

The initial selection of attribute sets may be performed by interview, intuition, or analysis. Protocol analysis is a subjective interview technique whereby the operator introspectively recounts the factors and procedures which enter into his decisions. Consciously considered attributes may be identified from this introspection. A second, more objective technique is possible if the decision situation is highly defined. The feature set can be determined according to the predictive error rate and from correlations of the candidate attributes (Felson, 1975). The first attribute chosen is that with the lowest expected probability of error (EPE). The EPE is the error rate that would result if the ith attribute alone were used as a basis for decision making. The second feature chosen is the one with the smallest correlation with the first attribute. The choice of the ith attribute depends on its correlation with the i-1 attributes already chosen.

Finally, the attributes in their raw form may be highly non-linear. Linear transformations to achieve interval scales are often warranted. The effect of the linear transformation is normally minor compared to the magnitudes of the test/retest reliability and intersubject differences (Edwards and Gutentag, 1975).

3.4 Adaptive Parameter Estimation

The previous sections have described the means of structuring the decision model, identifying relevant utility dimensions or attributes, and determining the levels of the attributes. Completion of this modeling process demands the assessment or inference of the subjective weights of each attribute.

It was noted in Section 2.3.5 that numerous techniques are available for off-line assessment of the operator's attribute weights. These techniques include direct elicitation of preference, decomposition

of complex decisions into hypothetical lotteries, and use of multi-variate methods to analyze binary preference expressions to determine underlying factors. These off-line techniques of utility assessment are accurate and reliable in many circumstances, but they have a number of disadvantages when applied to remote systems. Typically, these techniques require two separate stages--assessment and application. Assessment requires an interruption of the task and elicitation of responses to hypothetical decisions. Problems arise with such procedures since the operator's judgments may not transfer to the actual situation; the DM may not be able to accurately verbalize his preference structure (Macrimmon, 1973); and the judgments made in multidimensional situations are typically responses to non-generalizable extreme values (Keeney and Sicherman, 1975).

Estimation techniques relying on interference from in-task behavior may be more useful. These inference techniques assume a model of decision behavior and then fit the parameters of the model by observation and adjustment. The parameter estimation may be performed by multiple regression, time series analysis, heuristic search procedures, optimal control (Kalman filtering) techniques, pattern recognition, mathematical programming, or forms of iterative approximation. Each of these techniques has a specific domain of application. For remote system applications, two methods appear particularly useful. The multiple regression and pattern recognition techniques demonstrate the simplicity, robustness, and convergence guarantee necessary for on-line modeling and aiding.

The first of the two, multiple linear regression, is a highly efficient form of determining attribute weights from batches of behavioral observations. The technique uses a least-squared-error criterion to provide an unbiased estimate of the attribute weights. Confidence intervals on the estimated values may be determined at the same time. In fact, if factorial combinations of attribute levels can be presented to the decision

maker, a full analysis of variance can be made, providing data concerning the relative contributions of linear, quadratic, and higher-order terms (Macrimmon, 1973). Of course, both the regression and ANOVA techniques require large batches of data taken under comparable circumstances. Also, the responses given by the operator must be made along an interval scale. In-task choices between alternatives are not usable observations for regression. The regression techniques thus appear to primarily suited for separate, pre-mission simulations during which complex estimates are elicited. Once estimated, the parameter levels could be input to the remote system.

Actual in-task estimation appears feasible using pattern recognition techniques. Instead of batch processing, the pattern recognition methods refine the model decision-by-decision. Briefly, the technique considers the decision maker to respond to the characteristics of the various alternatives as patterns, classifying them according to preference. A linear discriminant function is used to predict this ordinal response behavior, and when amiss, is adjusted using error correcting procedures. This use of pattern recognition as a method for estimation of decision model parameters was apparently first suggested by Slagle. (1971). He made the key observation that the process of expected utility maximization involved a linear evaluation function that could be learned from a person's choices.

The suggested technique was soon applied by Freedy, Weisbrod, and Weltman (1973) to the modeling of decision behavior in a simulated intelligence gathering context. Freedy and his associates assumed the decision maker to maximize expected utility on each decision. They assigned a distinct utility, $U(x_{jk})$, to each possible combination of action and outcome, as shown in the decision tree in Figure 3-3. The probabilities of occurrence of each outcome j given each action k were determined using Bayesian techniques. These patterns of probability

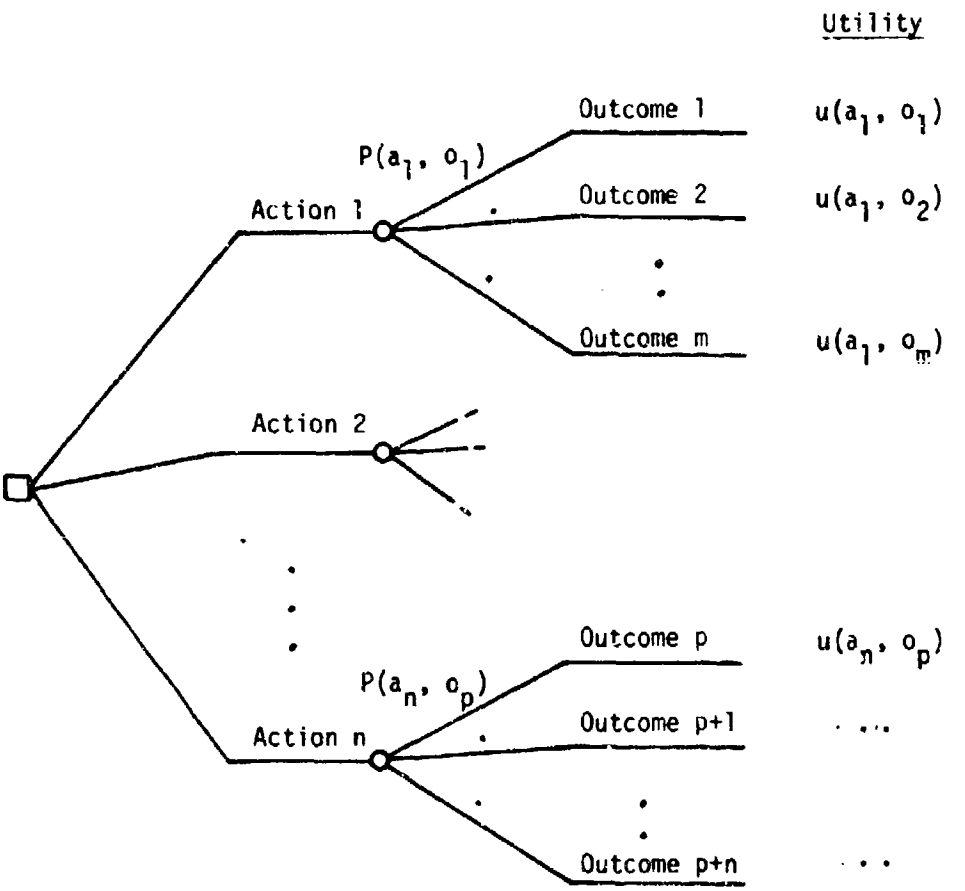

FIGURE 3-3. DECISION TREE OF UTILITY ESTIMATION PROGRAM DEVELOPED BY FREEDY ET. AL.

were used as inputs to the estimation program (Figure 3-3). The expected utility of each action A_k was then calculated by forming the dot product of the input probability vector and the respective utility vector. This operation is equivalent to the expected utility calculation:

$$EU(A_k) = \sum_j P(x_{jk}) \cdot U(x_{jk}) \tag{3-12}$$

The classification weight vector W_{jk} in the pattern recognition program acts as the utility $U(x_{jk})$. The alternative A_k having the maximum expected utility is selected by the model and compared with the decision maker's choice. If a discrepancy is observed an adjustment is made, as shown in Figure 3-4. The adjustment moves the utility vectors of the chosen and predicted actions (W_c and W_p, respectively) in the direction minimizing the prediction error. The adjustment consists of the following:

$$W_c' = W_c - d \cdot P_p \tag{3-15}$$

$$W_p' = W_p + d \cdot P_c \tag{3-16}$$

where

W_c' is the new vector of weights $[W(x_{1c}), W(X_{2c})]$ for action c

W_c is the previous weight vector for action c

d is the correction increment

P_i is the probability vector describing the distribution of outcomes $[P_{1k}, P_{2k}, \ldots P_{nk}]$ resulting from action k

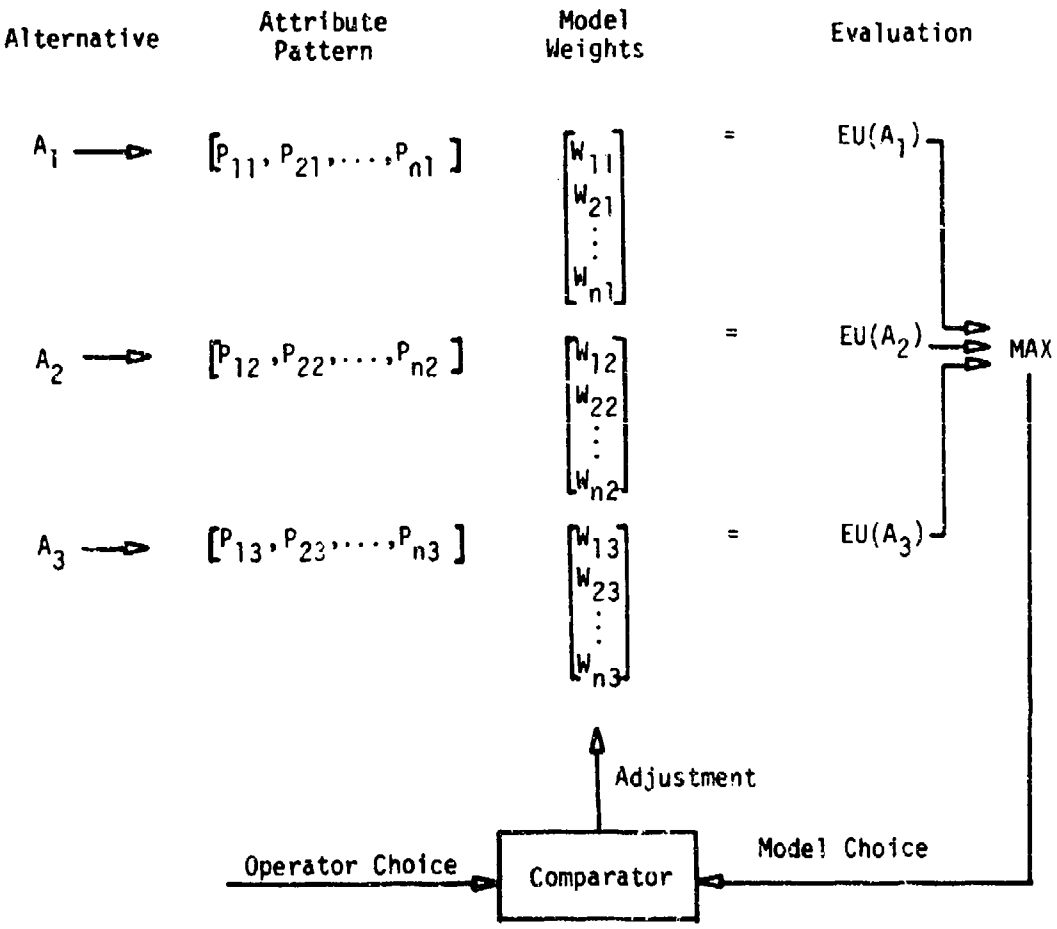

FIGURE 3-4. STRUCTURE OF UTILITY ESTIMATION PROGRAM OF FREEDY ET. AL.(1973)

This model is an adaptation of the R-category linear machine (Nilsson, 1965). The pattern classifier receives patterns of descriptive data (outcome probabilities) and responds with a decision to classify each of the patterns in one of R categories (actions). The classification is made on the basis of R linear discriminant functions, each of which corresponds to one of the R categories. The discriminant functions are of the form:

$$g_i(x) = W_i \cdot x \text{ for } i=1, 2, \ldots, R \qquad (3\text{-}17)$$

where x is the pattern vector and W_i is the weight vector. The pattern classifier computes the value of each discriminant function and selects the category i such that

$$g_i(x) > g_j(x) \qquad (3\text{-}18)$$

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

A geometric interpretation of the R-category linear machine is shown in Figure 3-5 (Nilsson, 1965). Decisions involving two possible consequences, x_1 and x_2, are evaluated according to three discriminant functions $G_1(\underline{x})$, $G_2(\underline{x})$, and $G_3(\underline{x})$. The lines of intersection between the discriminant hyperplanes are the points of indifference between actions. Mappings of these lines of intersection to the attribute plane are shown in the figure. The resulting regions R_1, R_2, and R_3 correspond to the actions maximizing the (expected utility) evaluation function.

The R-category technique becomes somewhat cumbersome if a large number of actions are possible or if the decision circumstances change rapidly. This problem is a result of the assignment of a distinct, holistic utility to each tip of the decision tree. The number of model

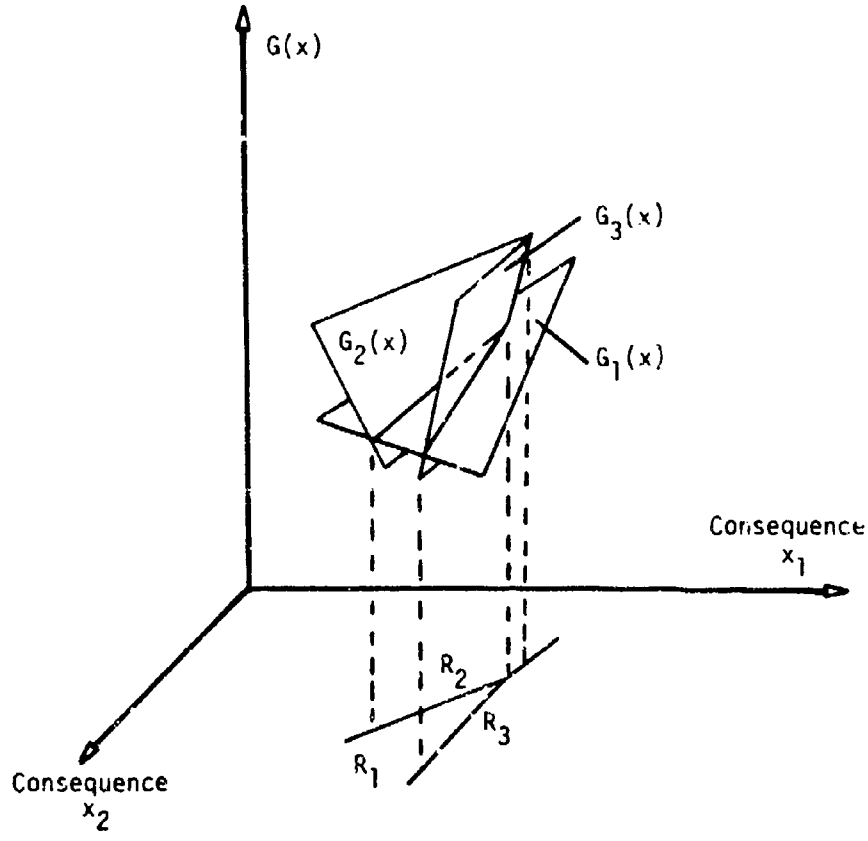

FIGURE 3-5. GEOMETRIC INTERPRETATION OF R-CATEGORY LINEAR MACHINE (ADAPTED FROM NILSSON, 1965)

parameters thus increases rapidly with an increase in the number of actions possible. Also, the only weight vectors adjusted in a given decision are those corresponding to the model-predicted and the actually chosen actions. This partial adjustment makes the system somewhat unresponsive to change.

Some of these shortcomings were lessened by a multi-attribute formulation developed independently by Felson (1975). Felson attempted to predict stock market behavior by fitting parameters of a linear model using pattern recognition techniques. Unlike Freedy, et al., Felson considered each action to be decomposable according to a single common set of attributes. Felson thus assumed that a single vector of weights could account for the observed behavior. The approach is centered around the use of a threshold logic unit (TLU), a two-category variant of the linear machine:

$$G(x) = \sum_{i=1}^{n} W_i x_i + W_0 \qquad (3\text{-}19)$$

where

W_i is the weight corresponding to attribute i
x_i is the level of attribute i
W_0 is a constant

Two possible consequences are considered in Felson's model--a rise or a fall in stock value compared to the market average. The consequence predicted depends on the sign of G(x). A single hyperplane serves to separate the two regions.

Figure 3-6 summarizes the estimation process of the stock market program. A set of feature or attribute levels (some of which are subjectively estimated) are input to the program. The attribute level

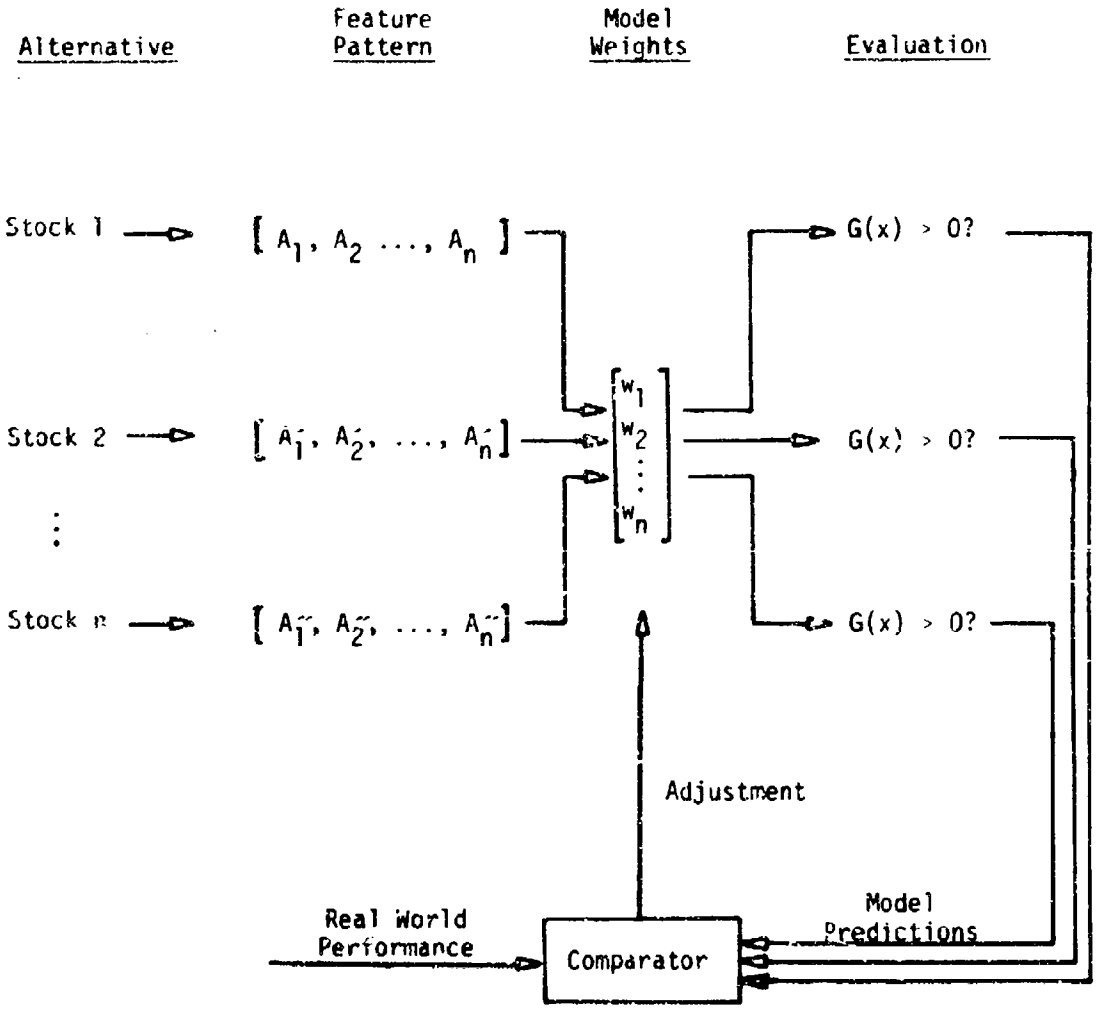

FIGURE 3-6. STRUCTURE OF FELSON'S (1975) MULTI-ATTRIBUTE STOCK PREDICTION PROGRAM

vectors are evaluated according to the current weight vector and a prediction of each alternative's market performance is made. The prediction is compared to actual stock performance (not to human behavior) and an adjustment made if a disagreement is present. The use of a single discriminant function adjusted at each erroneous decision led to very rapid training. Roughly 15 to 25 trials were found to result in asymptotic performance with 5 attributes (Felson, 1975). During this period, the error rate was found to drop from almost 50% to approximately 20%. Also, Felson noted the advantages of the pattern recognition approach over conventional estimation techniques: Its computational simplicity, its minimal need for initial information, and its parsimony of operation--change is made only when an error is detected.

A natural extension of Freedy's and Felson's approaches is to adapt the single discriminant, multi-attribute approach to the modeling of objective choice behavior. Each possible outcome of a decision can be associated with a set of attributes or objectives of the decision maker. An importance weight vector defined over the various attributes can then be adjusted to predict behavior. As in Felson's approach, the mechanism is simply that of a threshold logic unit. The adjustment rule following an incorrect prediction is

$$W' = W + d(x_c - x_p) \tag{3-20}$$

where

W' is the updated weighting vector
W is the previous weighting vector
x_p is the attribute pattern of the model-predicted choice
x_c is the attribute pattern of the decision maker's choice
d is the adjustment factor

The cycle of prediction, comparison, and adjustment of this proposed approach is illustrated in Figure 3-7. The model training is based on pairwise comparisons of alternatives, as shown in the right-hand portion of the figure. If a set of three or more alternatives is presented, and one is chosen, it is assumed that the DM prefers that alternative in any pairwise comparison with the remaining choices. Thus a single choice may result in a number of training adjustments.

The closed-loop nature of all of these programs is evident from Figure 3-8. It can be seen that the system compares the model output with the operator choice and uses the error as an input to the controller. Also, the system is adaptive in the sense of Gaines' (1972) criteria: the pattern classifier does not rely on a preset function to operate on the error, but it adjusts its parameters (the model weights) to minimize succeeding errors.

An immediate question concerning all of these models is whether the estimated parameters exhibit interval properties. The parameters are estimated solely from observations of ordinal responses. However, the resulting weight vector is defined along an interval scale. This is because in the limit, only a single hyperplane can correctly classify consistent responses to a fully represented pattern space. This hyperplane is also invarient to the usual positive linear transformations. The ensuing predictions made from these weights are not necessarily interval, though. The predictions are ordinal if SCUI is satisfied, and interval if both SCUI and marginality hold. These properties parallel those of the off-line technique of paired comparisons. The paired comparison judgments are ordinal; the resulting estimated parameters are interval; and the resulting predictions are again at least ordinal (Samet, 1976).

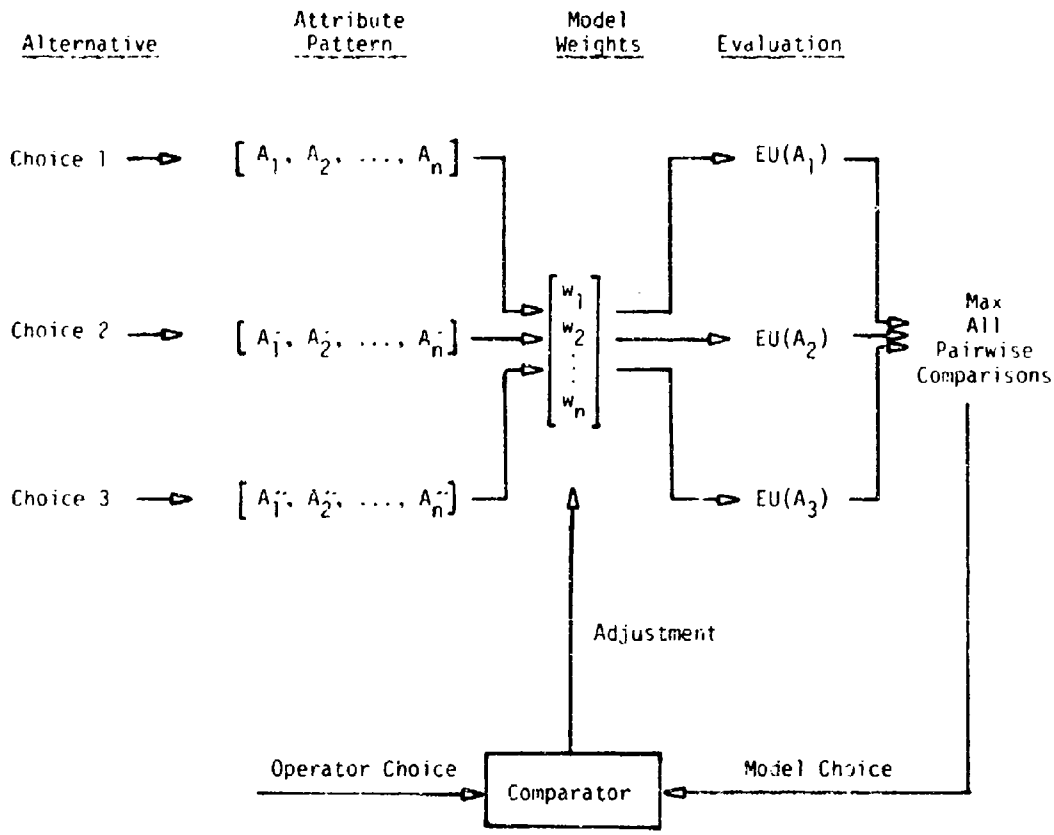

FIGURE 3-7. STRUCTURE OF PROPOSED MULTIATTRIBUTE ESTIMATION PROGRAM

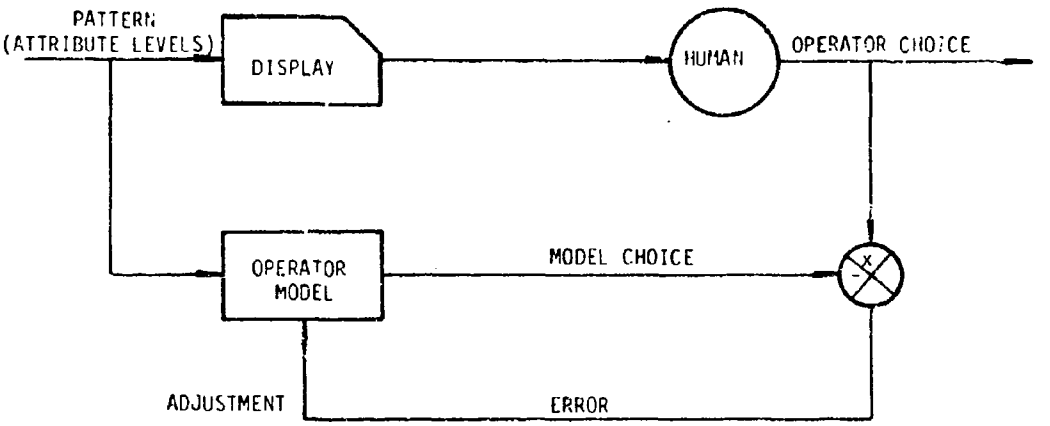

FIGURE 3-8. ADAPTIVE STRUCTURE OF ESTIMATION PROGRAMS

A possible advantage of the pattern recognition technique over many of the other forms of estimation is its flexibility of adjustment. Several types of error correction are possible for the TLU, each with a different combination of speed, stability, and complexity. The three principle forms are the fixed increment rule, the absolute correction rule, and the fractional correction rule. These differ solely in their formulation of the adjustment factor d in Equation 3-20.

The fixed increment rule simply assigns a non-zero constant to d. Thus the movement of the weight vector is a constant proportion of the difference in the predicted and chosen patterns. The correction may not be sufficient to avoid subsequent errors with the same pattern, but the process is eventually convergent (Duda and Hart, 1973). The fixed increment rule has the advantages of simplicity and relative insensitivity to inconsistent behavior.

A more rapid but also more potentially unstable rule is the absolute correction rule. This method sets d to be the smallest integer at which the error of the pattern is corrected. In the decision modeling situation, this becomes:

$$d = \text{smallest interger} > \frac{|k \cdot (x_c - x_p)|}{(x_c - x_p) \cdot (x_c - x_p)} \qquad (3\text{-}21)$$

in which

x_c is the attribute level vector of the operator selected choice

x_p is the attribute vector of the predicted choice

The fractional correction rule is similar to the absolute rule but is typically less extreme. The fractional rule moves the weight point some fraction of the above distance:

$$d = \frac{\lambda |k \cdot (x_c - x_p)|}{(x_c - x_p)(x_c - x_p)}$$

where λ is a constant $0 < \lambda < 2$.

All three of the adjustment rules have been proven convergent with linearly separable patterns (Nilsson, 1965). The speed of convergence is normally fastest with the absolute rule. This is illustrated for an example series of adjustments in Figure 3-9. The set of four numbered lines in the figure are a sequence of patterns. These patterns are shown as hyperplanes in a 2-dimensional weight space. Each hyperplane represents the difference between two multi-attribute vectors. The operator choice is shown by the direction of the arrow at each pattern. The absolute rule, (the triangles in the figure) is seen to achieve correct prediction after four observations, while the fixed rule (the circles) requires five. Unfortunately, the absolute rule is expected to be less forgiving of inconsistent behavior than the fixed or fractional rules. This is because of the large responses the absolute rule makes to operator inconsistencies. The fixed and fractional rules may exhibit a greater tendency to smooth or average the behavior.

3.5 The Adaptive Model as a Decision Aid

The adaptive decision model has the potential of improving system decision performance in two key areas:

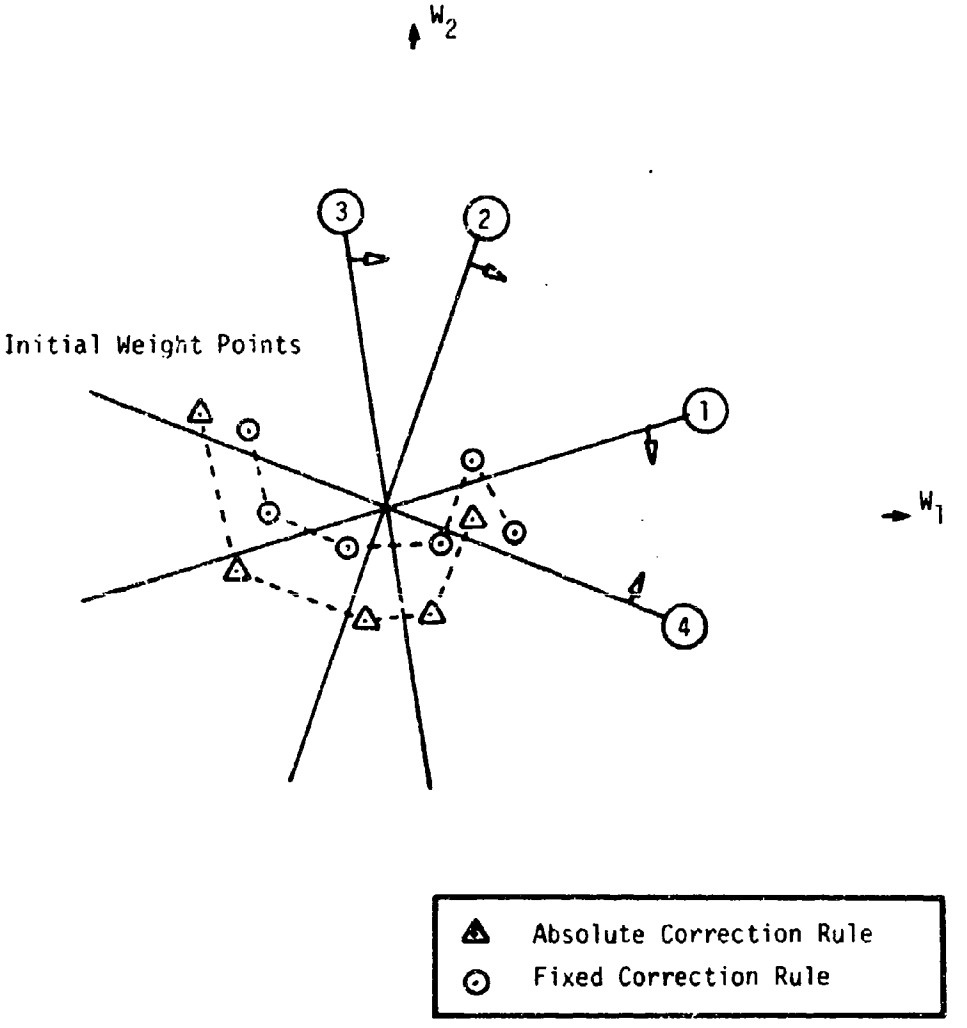

FIGURE 3-9. COMPARISON OF BEHAVIOR OF CONVERGENCE RULES

(1) Smoothing. The reduction of the random error or noise implicit in human response. This reduction is a consequence of the averaging of observations during parameter estimation.

(2) Augmentation. The amplification of the operator's decision making capability by displaying model recommendations. Observing the model recommendations, the operator may refine his behavior and possibly even consider a larger set of factors.

Smoothing or reduction of random effects in subjective weighting of data is a well-established advantage of linear models. Linear models based on an operator's average behavior typically outperform the actual behavior of the operator (Bowman, 1963; Goldberg, 1970; Dawes and Corrigan, 1974). Aiding by model recommendation of choices and by model-based automation should result in this type of performance enhancement.

The second area of improvement provided by the model, augmentation, deals with sub-optimal decision behavior that is more deep-seated than noise or random effects. Because of cognitive limitations, the operator can consider only a small number of attributes in a decision. In complex situations, he then constructs his own simplified and manageable model of the problem. This is Simon's (1957) "principle of bounded rationality" in which the man's behavior may be consistent with his own simplified model even though not even approximately optimal with respect to the real world.

The sub-optimal behavior resulting from cognitive limitations may possibly be reduced through model-based aiding. Macrimmon (1973) suggested that by operating in parallel with the DM, a model can present decision recommendations based on a normative processing of the circumstances and utilities. The operator's task is then changed to one of evaluation and

correction. Freedy and his associates (1976) displayed such model-based recommendations to operators in a simulated task of submarine surveillance. Significant imporvements in performance resulted, possibly from the opportunity to consider more complex and effective strategies.

Unfortunately, the parallel, closed loop relationship of man and model engenders some problems of dynamics. With aiding, the decision faced by the operator includes both the attribute patterns of the choices and a normative processing of those patterns. Since this processing is based on his previously observed behavior, it should lead to greater consistency, speed, and effectiveness in recurrent situations. However, it may result in inappropriate recommendations in completely new circumstances These characteristics are typical of predictive displays. The predictions are only accurate if future behavior can be estimated from previous observations. 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 the age. A recency bias of this type is realized to some extent by virtue of the adjustment mechanism. An additional bias may be necessary in rapidly changing situations.

The level of aiding provided by the model depends on the degree of training it has experienced. A possible sequence of training of the model, at least in the initial validation phase, is one of passive observation, then observation and recommendation, and finally, automation. At each succeeding stage, the model will gain more knowledge and become more independent. The first stage, observation, consists of the passive monitoring of decision conditions and operator choices. The initial arbitrary vector of attribute weights is adjusted and the model is sharpened with each incorrect model prediction. With experience, the

model approaches the behavior of the unaided operator. Once a minimum confidence or prediction level is reached, the model can aid the operator by making recommendations. In this second phase, the model provides a normative structuring for decisions, displaying the logical extrapolation of his previous behavior to the current choices. The model should, in time, stabilize to a consistent set of values reflecting the augmented decision strategy of the operator. Automation by the remote element can then begin. The automated decisions will still be subject to operator overrides, and the model parameters will continue to adapt, but the program will be largely autonomous. The model should then be capable of managing communications from the remote element.

4. SYSTEM APPLICATION

4.1 The RPV Communications Problem

4.1.1 Structure. The preceeding chapters have confronted the general problems of information evaluation and management in remote systems. It should be useful to explore these implications in a specific application. Probably the area of greatest immediate potential for machine control of communications is that of remotely piloted vehicle (RPV) supervision. Some operational RPV's already have an advanced degree of autonomy resident in their autopilot systems. At the same time, the evaluation and goal direction functions are the responsibility of the human operator, so that some degree of interaction is essential. Finally, substantial communication costs are present because of the possibility of detection, the energies expended in data transmission, and the supervisory loads imposed on the operator. The three elements above - machine intelligence, operator supervisory requirements, and costly communications - are factors that encourage the placement of the communications evaluation and management functions with the remote element.

The RPV control task is normally hierarchical and goal directed in nature. The levels of function range from continuous stabilization adjustments to long range planning of the overall route. Figure 4-1 (adapted from Roscoe and Eisele, 1976) depicts a representative ordering of these functions along with the feedback loops involved. Usually the lower level (high frequency) functions such as vehicle stabilization are automated. For these vehicle control functions, the complete processes of actuation, performance measurement, and comparison with objectives are performed by the autopilot system.

The intermediate level functions, on the other hand, tend to be actions assumable by either the remote element or the human. These

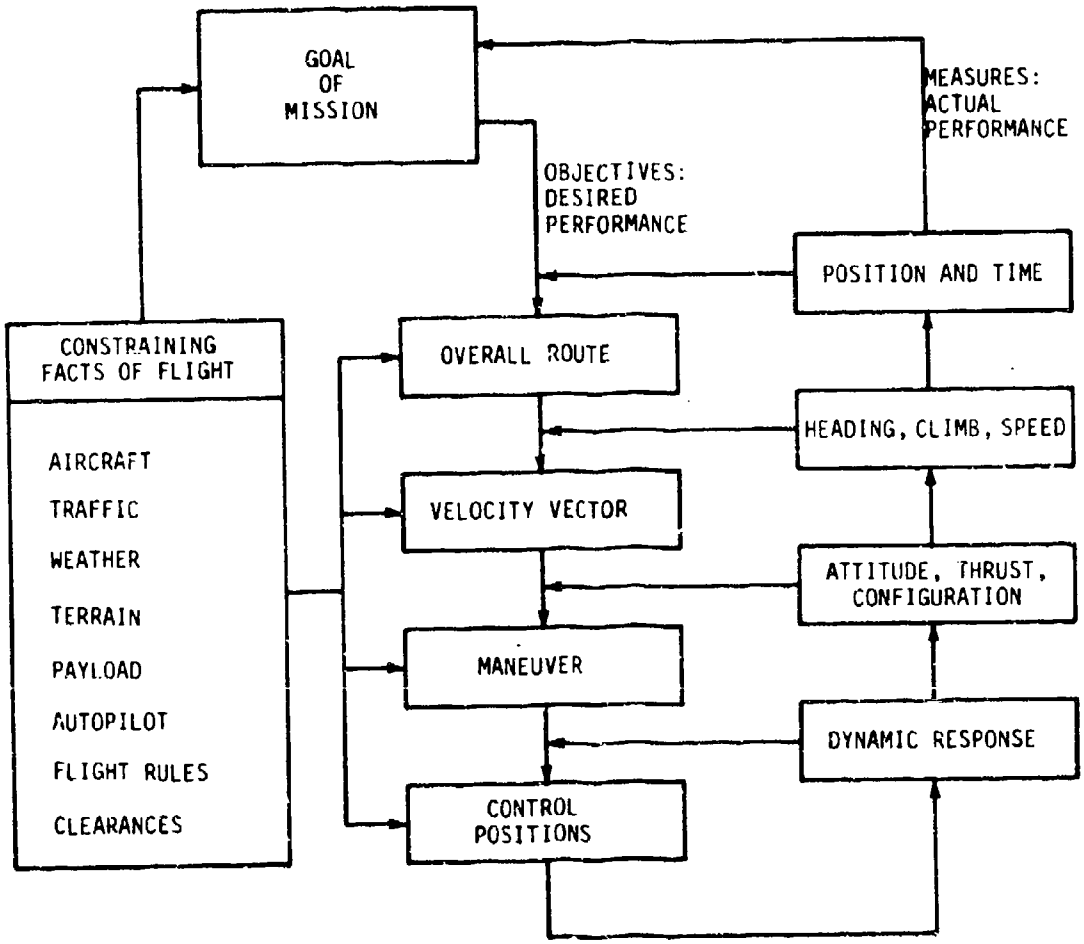

FIGURE 4-1. HIERARCHICAL STRUCTURE OF RPV CONTROL TASK

functions involve such actions as responses to unforseen obstacles, identification of targets, and handling of system failures. Such actions tend to be discrete decisions rather than continuous control responses. Speed of response and minimization of costs dictate that the machine take responsibility in some cases, while flexibility and safety necessitate human control in other situations (Johansson, 1976).

The highest level functions, such as definition of the system objectives and constraints, are almost exclusively the domain of the human operator. For example, the criteria of performance at each level of the functional hierarchy are input by the human. Certain constraints, such as traffic, weather and terrain conditions, may be recognizable by the machine element, but the overall decision policy is virtually always defined by the human.

4.1.2 RPV Mission Characteristics. The degree of functional responsibility assumable by the machine depends to a large extent on the task circumstances. Typically, a remote vehicle mission is defined by a series of mission phases. The phases can be characterized by the amount of communications allowed, the availability of feedback concerning vehicle and environmental states, the probability and extent of potential losses, and the time available for decision making (Mills, Bachert, and Rume, 1975). Each of these factors influence the degree of autonomy that can be realized by a remote system.

The amount of communications allowed is a function of the channel capacity, the direct and indirect costs of transmission, and the amount of attention the human operator can contribute. Channel capacity is defined by such factors as band width, time delay, and signal-to-noise ratio. These factors become more troublesome as the distance to the remote element becomes greater and as the number of intervening obstacles increase. The direct costs of communication also increase with distance and number of hazards. These costs include energy expenditures and equipment expenses.

The indirect costs - increased possibilities of detection, countermeasures, etc. - are more a function of the hazardousness of the region rather than the communications distance. The available operator attention, finally, is defined by the number of controlled systems, the secondary task demands, and the individual capabilities of the operators.

The costs and payoffs associated with the various possible control outcomes also vary with mission phase. The consequences are defined not only in terms of attrition of equipment and attainment of objectives, but also as a function of organizational policy and procedures. The relative importance of fuel expenditures, vehicle survival, countermeasures, etc., change as the mission objective is approached, attained, or past (Fogel, Englund, Mout, and Hertz, 1974). The relative importance of these factors must be assigned by the human operator or by the organization.

Available time for decision making varies throughout the RPV mission as a direct function of the varying vehicle speed, altitude, and surrounding weather conditions. Altitude and weather determine the distance that obstacles, navigation points, or targets can be observed. The speed then determines the available time. Decision time can be expected to influence the amount of information that can be processed and the probability distribution of the possible consequences.

In sum, the selection of information and control to allocate to a supervisory human operator is a complex and dynamic decision. The decision maker must continually weigh the probable usefulness of the information or assuming control -- energy costs, attention requirements, risk of detection, etc. These judgments often must be based on subjective factors, as the decision task is normally too complex and dynamic to be analytically tractable.

4.1.3 System Overview. An overview of the sequence of processes involved in RPV supervision is diagrammed in Figure 4.2. The sequence is divided into three segments - modeling, analysis, and execution. Modeling is considered to consist of structuring and assessment. Structuring is the definition of the various components of the decision model, while assessment is the determination of the parameter levels. The modeling segment is shared in function, as the human operator typically defines the decision structuring (at least until self-organizing systems can be realized) and the computer performs the assessment. The second segment in the cycle, analysis, is assigned completely to the computer. Analysis involves solving a model to determine its implications. Analysis also involves computing the effects of altering model assumptions. The final segment, execution, is again a flexible function: either man or machine may make the decision. In the early stages of model training, the human would be executed to perform the action with the machine observing passively. Later, with increased confidence in its controls, the machine could either make recommendations to the operator or take over the decision function (subject to operator override). The coming sections will consider in greater detail the stages of modeling, analysis, and execution.

4.2 Modeling

4.2.1 General. The multi-attribute model developed initially in Section 3.4 will provide the basis for the structuring and assessment processes. The specific steps of these modeling processes are outlined in Figure 4-3 through 4-5 (adapted from Gardiner, 1975). The first figure shows the two sides of the modeling problem, probability estimation and utility assessment. The upper portion of Figure 4-3 details the processes of probability estimation. These include delineation of the possible states of the environment, evaluating the current level of uncertainty concerning the states, selecting information to reduce the uncertainty, and revising the probability estimates in light of the new data.

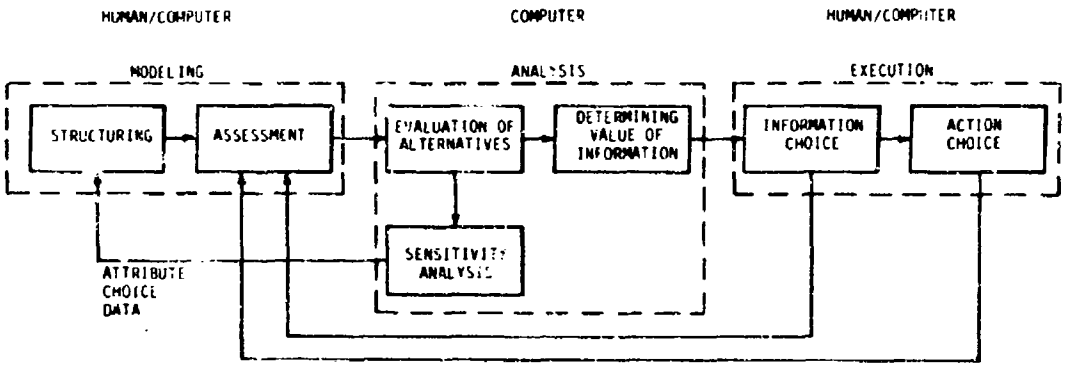

FIGURE 4-2. THE PROCESS OF MODELING, ANALYZING, AND EXECUTING DECISIONS

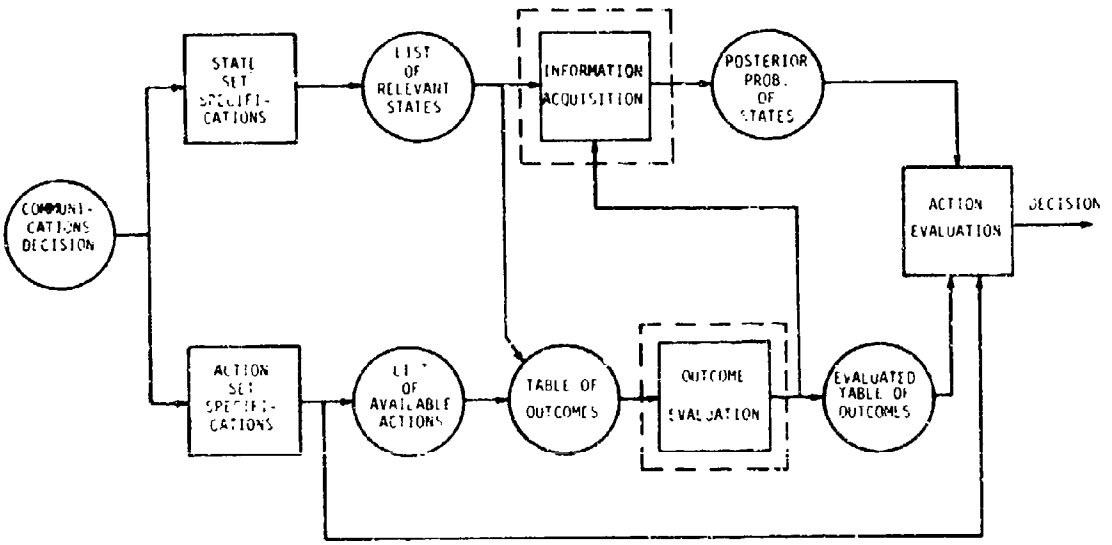

FIGURE 4-3. DECISION PROCESS CHART
(Adapted from Gardiner, 1975)

The key element in this probability estimation sequence is the information acquisition stage (enclosed by dotted lines). Figure 4-4 elaborates this stage, showing the steps that go into the choice of information and the subsequent incorporation of the datum into the situation estimate. The upper portion of the figure deals with the information source selection. The characteristics of the various available sources are determined by observation and analysis. This estimation of the characteristics of the information sources is accomplished by successive comparisons of messages received and subsequently observed states. The choice of information source is then made according to the potential impact of the information on the prior probability estimate. Once a source is selected and a datum observed, the information is incorporated into a revised situation estimate through Bayes' rule (see Equation 2-2).

The other major modeling process is utility assessment or outcome evaluation. The possible combinations of actions and states are enumerated off-line prior to a mission. The problem is then to assign consequence levels and importance weights along a predefined set of dimensions. Figure 4-5 elaborates this process. The first step is the selection of an independent, exhaustive, and predictive attribute set. The attributes are the various constituent aspects of the consequences. Each combination of action and outcome is associated with a set of attribute levels. This is done by observation and adjustment, just as in the determination of information source characteristics. Scaling procedures are applied to the raw consequence dimensions to arrive at normalized values. Each attribute is scaled so that its plausible range spans zero to one.

The attribute weight estimations (the W_i in Equation 3-6) can be performed using either pattern recognition techniques or off-line decomposition approaches. These methods were described previously in Section 3.4 and will be applied to the RPV control situation in the coming sections.

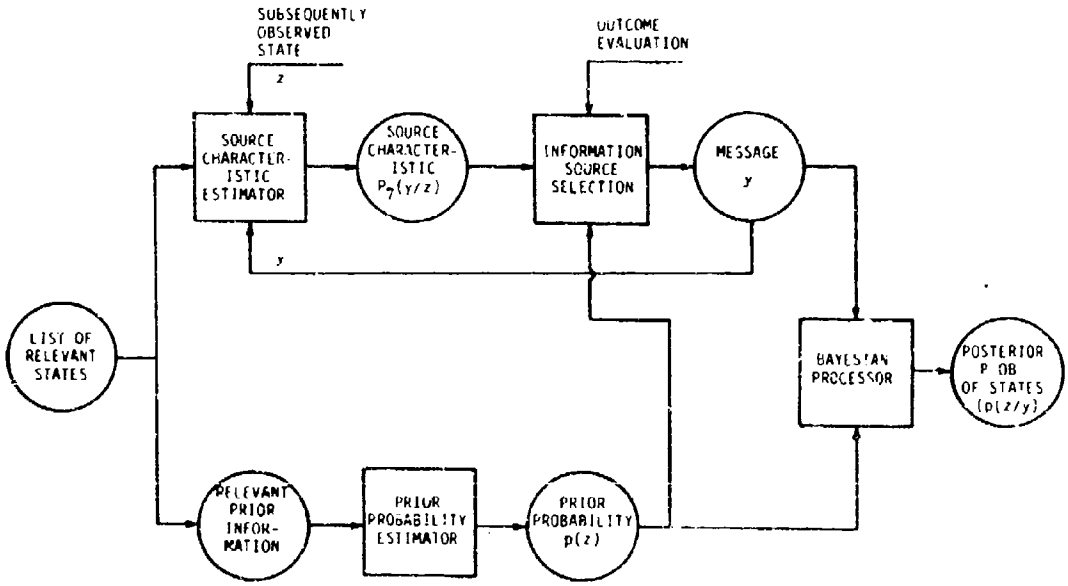

FIGURE 4-4. PROCESSES INVOLVED IN PROBABILITY ESTIMATION

(This is an elaboration of the "Information Acquisition" block of Figure 4-3.)

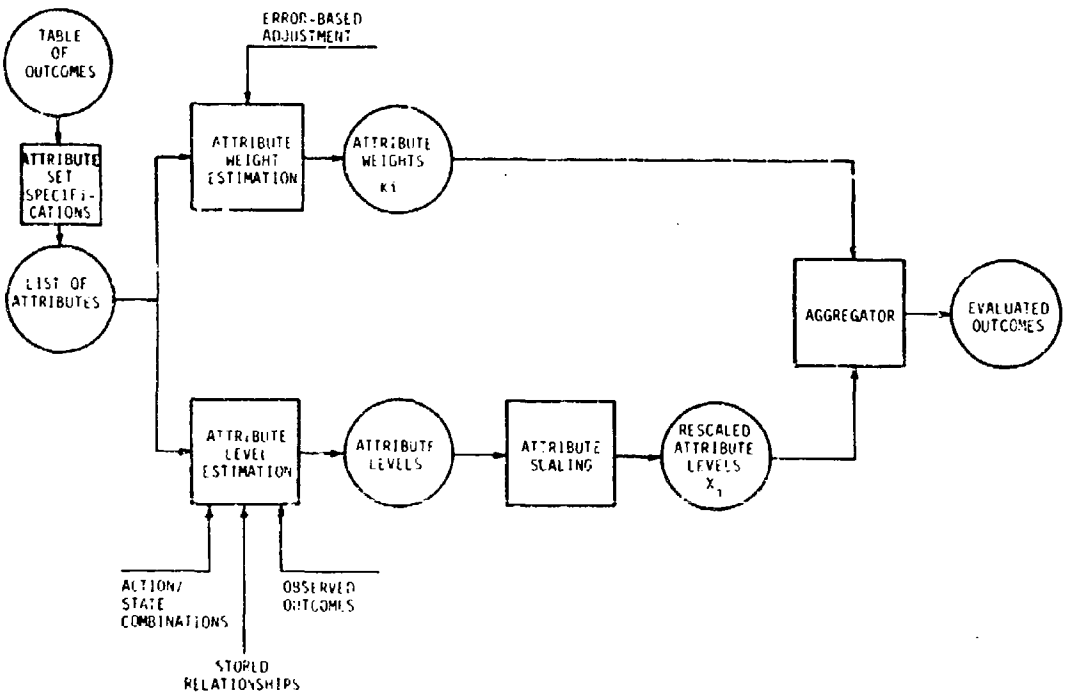

FIGURE 4-5. PROCESSES INVOLVED IN OUTCOME EVALUATION

(This is an elaboration of the "Outcome Evaluation" block of Figure 4-3.)

The above process description is, of course, still a simplification of the general information seeking problem. Certain aspects cannot be considered at this time. For example, continued sampling of information prior to the action decision is not represented. A complication of special note is the dynamic nature of the environment. The effects of an action would be expected to be different depending on the time of execution. Timing, however, may not be a major problem in remote systems applications, since the action choices are normally made at times determined by the situation. Typically, the decision is forced by the sensing of an obstacle or critical situation. Thus, the consequence set does not necessarily have to be time dependent.

The following sections will develop some of the specific points of the modeling cycle.

4.2.2 Probability Estimation. The two major probability parameters requiring estimation are the prior probabilities ($P(z)$) and the conditional probabilities $P(y|z)$. The priors are presumed to be of only minor importance in remote systems applications. This is because highly diagnostic data from the remote system sensors should result in virtually the same posterior probability estimates regardless of the values of the priors. Diagnosticity refers to the informativeness of the data concerning the states z_i. A highly diagnostic datum y exhibits a high likelihood ratio:

$$L_{12} = \frac{P(y|z_1)}{P(y|z_2)} \tag{4-1}$$

When incorporated into Bayes' rule, such a datum will have a major effect regardless of the prior probabilities of z_1 and z_2. It is expected that a remote system with sophisticated sensors operating in a well defined environment will receive data of high diagnosticity.

This reduction of importance of the priors is fortunate, as estimates of $P(z)$ can be only coarsely estimated prior to an operation or mission.

The priors are descriptors of the mission phases -- estimates of the likelihood of weather problems, adversaries, terrain obstacles, etc. These estimates are by nature of low confidence.

The sensor characteristics $P(y|z)$ are easier to estimate accurately. This is because the sensor characteristics are assumed to be invarient over time, unlike the changing prior probabilities. Comparisons of the messages received and the states subsequently observed provide the necessary data. $P(y|z)$ can then be derived from frequency counts of $P(z)$ and $P(z|y)$ using the following expression:

$$P(y|z) = P(z|y)\ P(y) \qquad (4\text{-}2)$$

These observations may be made either in a simulation or during actual system operation.

4.2.3 Factor Choice. It was noted in Section 3.3.4 that the attribute set should be accessible, monotonic, independent, complete and meaningful. Also, a single set must account for both information acquisition and action selection behavior. Finally, the attribute set must be manageably small in dimension. With these considerations in mind, an initial taxonomy of consequences can be organized around the following five areas:

(1) Communications Costs - The expenditures associated with use of the communication channel. These may include requirements of energy, equipment, and operator attention.

(2) Equipment Attrition - The consequences of control concerning the integrity of the vehicle. Included are fuel expenditures, system damage, and vehicle loss.

(3) Objective Attainment - The degree of accomplishment of the mission objectives. Target goals may be the area reconnoitered, payload delivered, and political impact obtained.

(4) Dynamic Effects - The future consequences resulting from the current actions. These consequences may include effects on subsequent autopilot capabilities, availability of future information, and changes in the environment resulting from the action.

(5) Subjective Needs - The operator may have propensities for obtaining (or refusing) information or for maintaining control beyond that called for by the above factors. These preferences reflect needs of task continuity, maintenance of load, or other idiosyncratic factors.

A useful consequence set might contain a single dimension or attribute from each of these categories. In fact, five attributes appears to be an upper limit to the number of factors a decision maker can effectively consider (V. Winterfeldt, 1975). If several factors contribute to one consequence dimension, these factors should be combined using a single common scale -- dollars, ship-equivalents, fuel quantity, etc.

Each of the attributes -- communications costs, vehicle losses, etc. -- must be scaled with interval properties along a set range. The least considerable consequence that may occur is assigned a level of zero on the scale. The most desirable consequence is assigned a level of one. The weighting factors W_i should also be normalized so that the overall worst combination of factors results in a value of zero and the overall best combination a value of one. The method of assessment of these weights will be discussed shortly.

Probabilistic consequences will be computed according to their expected value. For example, the vehicle loss attribute may have three possible levels, each with a different estimated probability of occurrence. The expected value is computed by the following additive expression:

$$E(x_i) = \sum_k P(z_k) x_{ijk} \qquad (4\text{-}3)$$

where $E(x_i)$ is the expected consequence level

P_k is the probability of state k occurring

x_{ijk} is the level of attribute associated with action j and state k

4.2.4 Consequence Level Determination. The actual level of each of the attributes for a given outcome can be determined by mappings between predictive features and the attributes. Predictive features must be identified which are accessible to an onboard program and capable of determining the consequence levels. Mappings between the predictive features and the attributes are either pre-established or determined by observation and adjustment.

The data available to the decision program are:

(1) Directly sensed information concerning the environmental state (weather, terrain, adversaries).

(2) The vehicle state (velocity, fuel, autopilot capability).

(3) The channel characteristics (capacity, noise, cost).

(4) Operator capabilities (attention, load).

(5) Communications choices (information, control acquired).

A manageable subject of these features must be determined. This can be done using the correlational procedure described in Section 3.3.5. The consequence mapping can then be refined by comparison of the predicted and actually observed consequences, as in Figure 4-6. The mapping can be developed either by prior definition, by regression, or by the pattern recogniti techniques described in the coming section.

4.2.5 Weight Assessment. The method of assessment developed in Section 3.4 -- adaptive estimation using pattern recognition -- appears well suited to the remote system problem. The goal is to estimate the operator's decision making policy by observation of his choices. The procedure is diagramed in Figure 4-7. First, expected consequence vectors associated with each combination of information and control are input to the model. These consequence vectors are dotted with the weight vector, resulting in evaluations along a single scale. The maximum expected utility choice is determined and compared with the operator's actual choice. If a discrepancy occurs, the weight vector is adjusted according to the procedures outlined in Section 3.4. Ideally, the error correction moves the weight vector in a direction minimizing subsequent errors.

The criteria used for model training is a question of special importance. The model training could be based on objective (external) criteria or on subjective (internal) criteria. It was noted previously that external criteria of system performance, such as control error, speed, and monetary costs are seldom available during the execution of a mission. Felson's (1975a) stock market program, adjusted according to daily stock performance, is a notable exception. More frequently, a completely analytical model is impractical and subjective criteria for model training must be used. Here it is assumed that the operator's behavior reflects the task objectives along with individual needs. The adaptive model functions in parallel with the operator, attempting to capture his decision policy by fitting a normative framework to his choices. This was shown earlier in Figure 4-2 by the feedback from the decision execution stage to the assessment block.

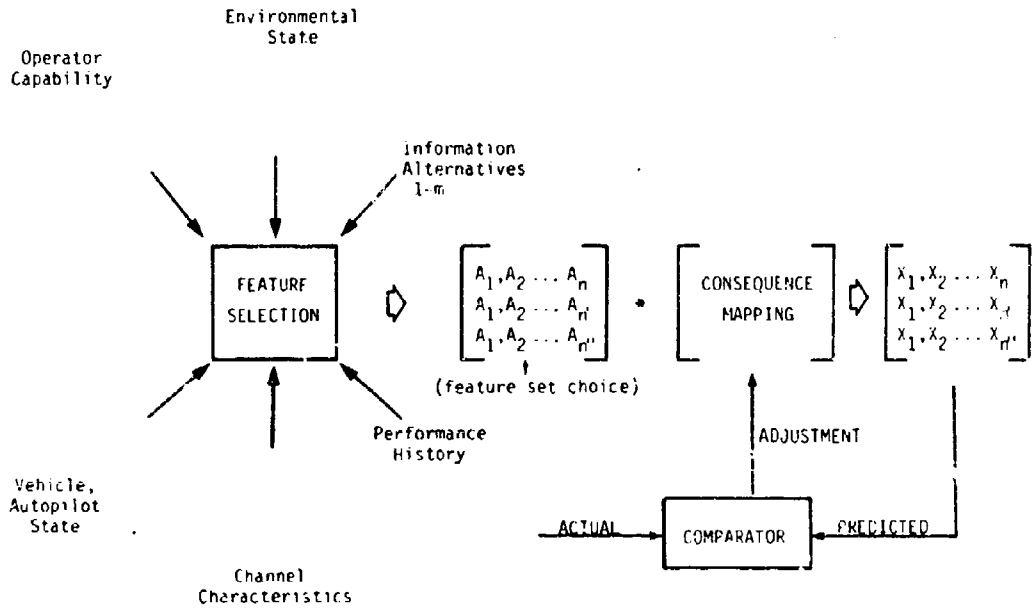

FIGURE 4-6. ADAPTIVE METHOD FOR ESTIMATION OF CONSEQUENCE MAPPING

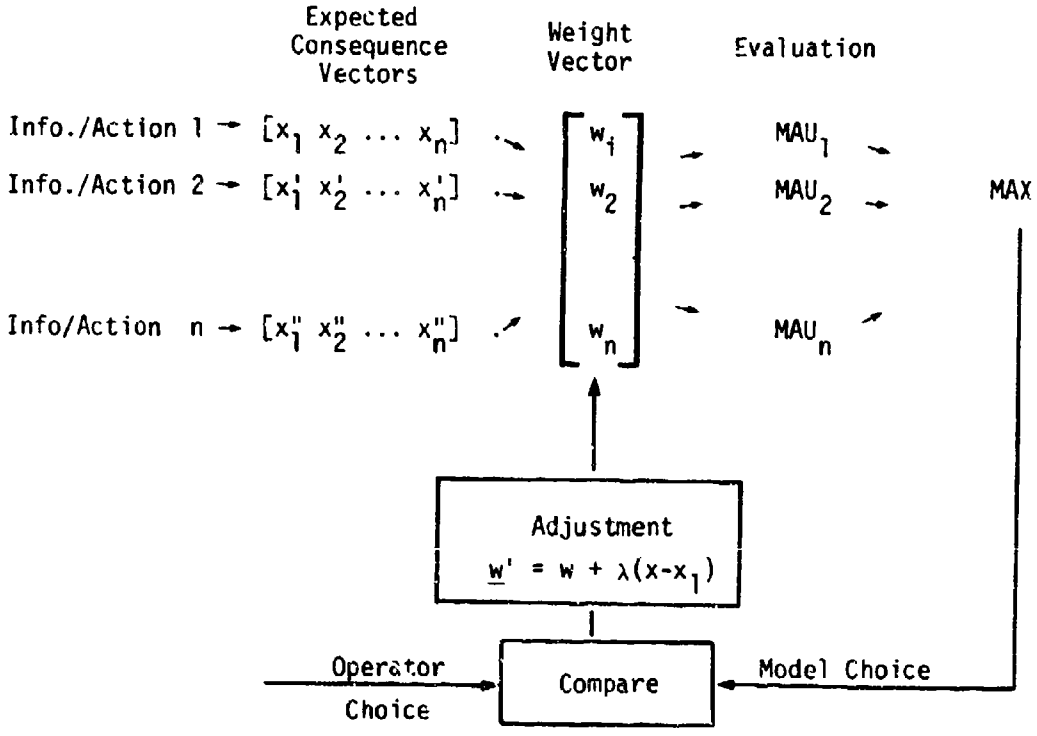

FIGURE 4-7. ADAPTIVE ESTIMATION PROCESS

In many situations, an occasional indicator of objective performance is observable. The RPV may be lost, the target attained, or other goals accomplished. In this way, the correctness of a sequence of subjective decisions may become known. The utility model would still be trained by observation of the operator's choices. If the sequence of choices led to an objectively favorable outcome, the new parameter set would be retained. If the outcome was unfavorable, the parameter set would be returned to the levels present prior to the sequence of decisions. In this way, objective criteria would guide training, but the explicit decision-by-decision policy for controlling the RPV would be subjectively derived.

Of course, the adaptive techniques of estimation described above are warranted only if repetitive decisions are available for training and if the weight differences present are important. In cases where only a few decisions will be made, off-line estimates of the weights W_i are favored. Here, techniques such as direct estimates, hypothetical lotteries, or paired comparisons are used for estimation prior to the mission. Some problems may occur since it is assumed with these techniques that the system requirements will not change after the estimates. These techniques also assume that the operator can effectively express his preferences along each dimension of choice.

Questions concerning the importance of weight differences are more basic. It was noted in Section 2.3.4 that unit weighting schemes (in which all weights W_i are set equal to 1.0) can be quite effective in certain circumstances. Errors in the model form, positive correlations between variables, and small sample sizes all reduce the predictive capabilities of inferred weights compared to unit weights. Essentially, the more precise and parsimonious the model, the more important inferred weights are.

Unit weighting schemes are expected to see only minor application in remote systems modeling. Careful selection of attributes minimizes

intercorrelations between variables, and the correlations that do occur should tend to be negative. For example, in RPV supervision, costly information is generally more informative than inexpensive information, and equipment attrition tends to be negatively correlated with goal attainment. These circumstances favor inferred weight models. The unit weighting schemes should primarily be useful as starting points for estimation, or as strategies for situations in which a great deal of noise is present.

4.3 Analysis

The analysis functions are computational processes intended to determine the model implications and sensitivity. Analysis includes such processes as evaluation of the various information sources and specification of the types of information needed. Also, sensitivity analysis may be made regarding changes introduced in various aspects of the model.

The type of information needed in a particular situation can be specified analytically by working through the predictive features and consequences, or empirically by sensitivity analysis. The analytical approach requires that relationships between the type of information and the consequences can be specified. Then if situational requirements result in certain consequences being emphasized, the corresponding forms of information can be identified.

The empirical approach utilizing sensitivity analysis is probably the most practical means of developing design criteria and determining model characteristics. Sensitivity analysis involves the systematic alteration of input variables to see how such changes affect outcome variables. The parameters that can be varied are:

(1) Information Sources - Source characteristics, costs.
(2) Situational Factors - Prior probabilities, predictive features.

(3) Consequence Levels - Values of each consequence dimension.
(4) Importance Weights - The inferred or preset W_i.

The possible criterion or output variables are also numerous:

(1) Predictive Capability - Percent of decisions predicted.
(2) Speed of Convergence - Number of decisions or time required for training.
(3) Objective Performance - Level of task performance.

Sensitivity tests also disclose whether a flat maxima situation is present. Here, large changes in an input variable lead to only minor changes in the output. It was noted earlier that linear models are often insentitive to differences in attribute weights -- witness the efficiency of arbitrary unit weightings in many situations. However, Slovik, Fischoff and Lichtenstein (1977) concluded that this flat maxima behavior is primarily a problem of continuous choices. With discrete choices, (e.g., perform surgery vs. don't perform surgery) it has been shown that a moderate error in probability estimation can lead to a substantial decrease in expected utility. This is expected to be the situation in RPV supervision.

4.4 Preliminary Tests of the Adaptive Model

Prior to experimentation with human subjects, a number of automated simulations were run testing the pattern recognition model. Human choice behavior was simulated by programming a "model operator" to make decisions according to pre-set attribute weights. Sets of consequence vectors (attribute levels) were generated randomly and presented to the model operator as pairwise choices. The model operator selected the choice determined by the preset weights, while a separate adaptive program observed the choices and modeled the behavior.

The factors tested with this simulation were the number of consequence dimensions, the type of adjustment rule, and the degree of consistency of the operator. The adjustment rules tested were the fixed and the absolute rules (see Section 3.4 for a description of these rules). The operator consistency was controlled by adding random numbers to the pre-set weights on a set percentage of decisions.

The model was found to converge at a rate dependent on the number of attributes and the degree of operator consistency. Only minor differences in dynamics were seen between the fixed and absolute adjustment rules. Convergence behavior as a function of the number of attributes is shown in Figure 4-8. The number of pairwise choices necessary to achieve a criterion level of prediction increases rapidly with dimensionality. As little as 15 decisions resulted in asymptotic training with three attributes, while 75 decisions were required for seven attributes. Figures 4-9 and 4-10 show this behavior more clearly for representative cases of 3 and 5 attributes. For three attributes, the adaptive model was able to rank the weights correctly after 12 decisions and achieved 100 percent prediction (for a twenty-decision window) after 19 decisions. The corresponding values for 5 attributes were 28 and 50 decisions, respectively.

The introduction of inconsistent behavior also increased the required training time. Figure 4-11 shows the increase in training time due to the imposition of 20 percent inconsistent behavior on the model operator. The number of decisions to convergence roughly doubled with this level of inconsistency.

The number of decisions required for training in an RPV supervision task should be considerably less than the number demonstrated in this simulation. The operator normally selects an information and control choice from a sizable number of alternatives rather than from a single pair. The

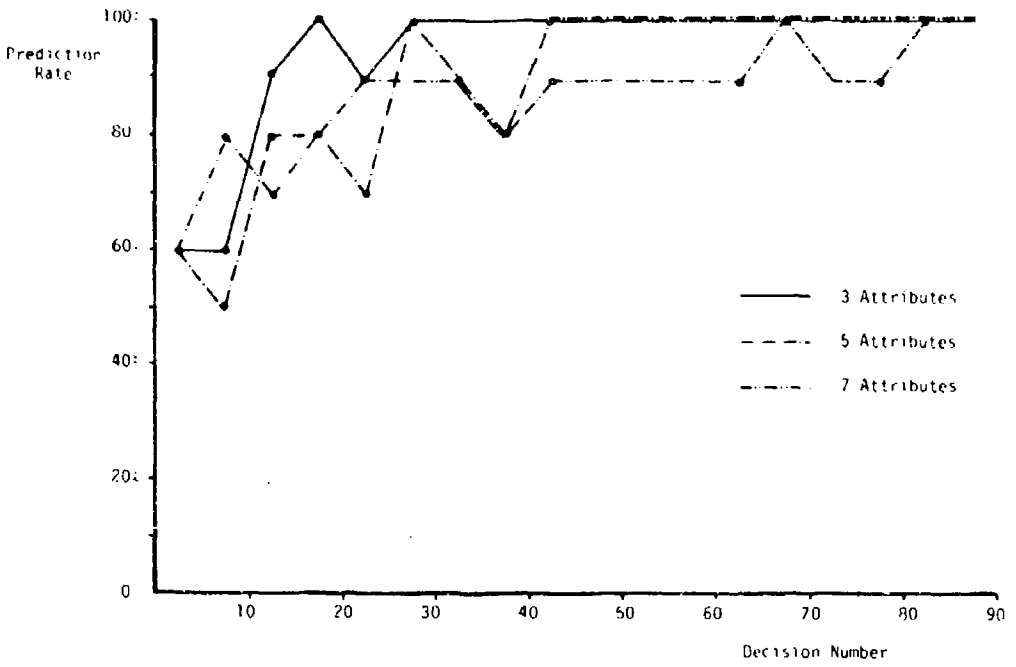

FIGURE 4-8. CONVERGENCE BEHAVIOR OF THE ADAPTIVE MODEL AS A FUNCTION OF THE NUMBER OF ATTRIBUTES

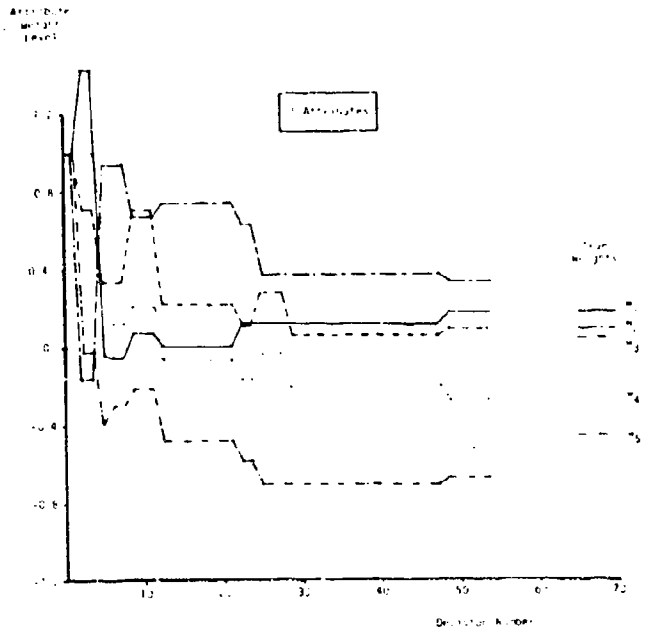

FIGURE 4-9. STRATEGY ACQUISITION FOR 3 ATTRIBUTES

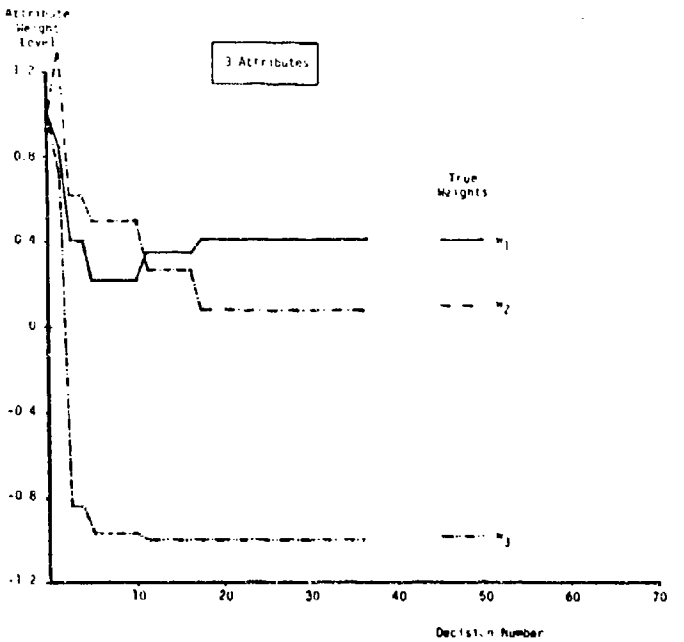

FIGURE 4-10. STRATEGY ACQUISITION FOR 5 ATTRIBUTES

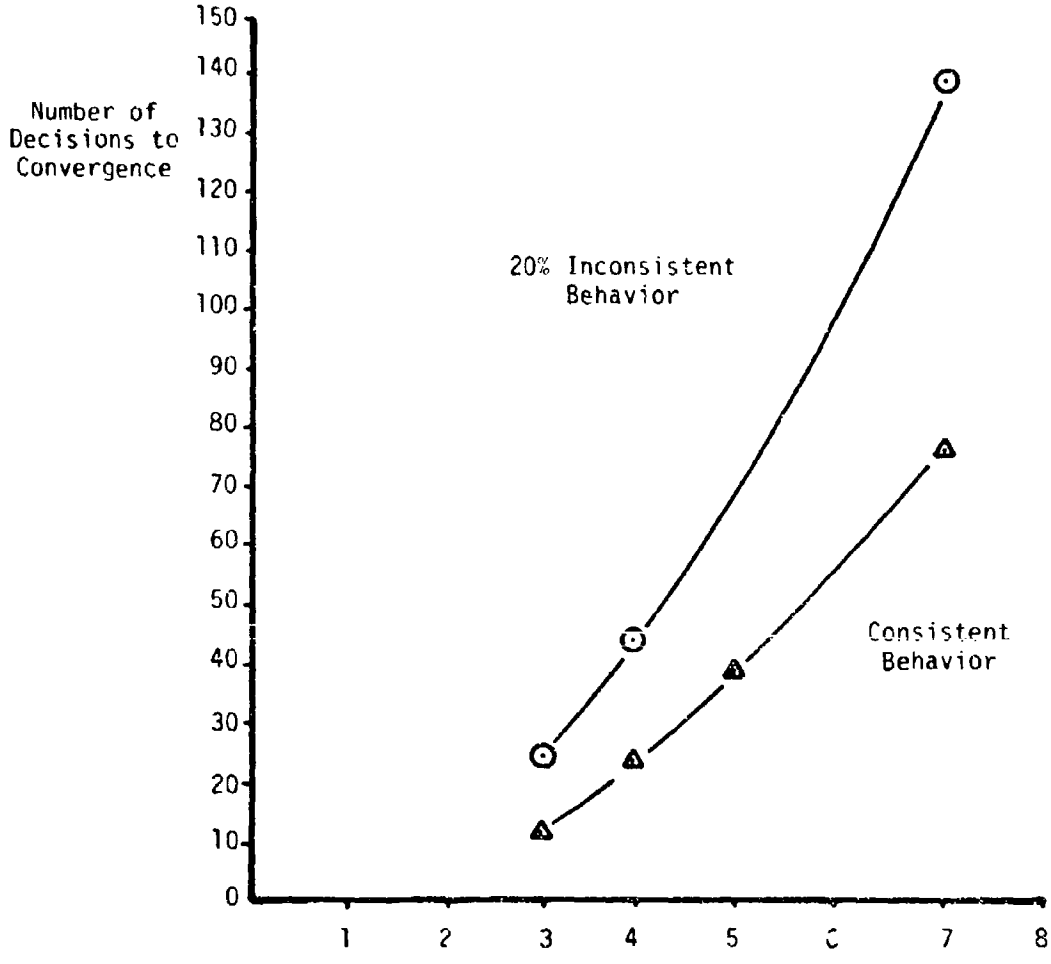

FIGURE 4-11. INFLUENCE OF OPERATOR INCONSISTENCY ON CONVERGENCE PERIOD

frequency of adjustment should thus increase, accelerating the training. Also, the model produces a ranking of the various alternatives. If the operator selects an alternative that lies in the lower range of the ranking, pairwise comparisons and adjustments can be made with all alternatives ranked above the operator's choice. This further reduces the training period.

5. EXPERIMENTAL STUDY

5.1 Overview

An exploratory study was performed to test the effectiveness of the of the adaptive multi-attribute model with human subjects. The study utilized a task simulation resembling control of a remotely piloted vehicle (RPV). Individual subjects were required to navigate the RPV through a changing, hazardous environment. In doing so, the operators were able to select different combinations of information display and control allocation. The main objective of the study was to determine the ability of the decision model to analyze, predict, and aid in these information and control choices.

5.2 Hypotheses

The following experimental hypotheses were tested:

(1) The adaptive model can accurately predict operator information and control choices under a variety of task conditions.

(2) The model-estimated parameters are more predictive and demonstrate greater construct validity than a unity weighting scheme (an arbitrary model with all weights set to 1.0).

(3) Subjective biases and inconsistencies of the operator can be identified using the adaptive model.

(4) Aiding provided through display of the model recommendations will result in performance superior to that obtained: (1) without aiding; and (2) with aiding derived from a unity weight model.

5.3 Task Simulation

The task simulation was patterned after an important and representative remote system task -- communication with and control of an RPV. This simulation is an adaptation of an RPV supervision task developed previously for study of human factors aspects of shared decision making (Steeb, Artof, Crooks, and Weltman, 1976). This task appears appropriate as it combines some degree of fidelity with an extensive amount of experimental control. Briefly, the updated simulation requires the operator to supervise control of an RPV in a hazardous environment. The environment contains obstacles of uncertain form and extent. At each set of obstacles, the operator can either control the RPV manually, or he can delegate control to the onboard autopilot system. Also, the operator has the option of accessing several forms of information about the environment. The forms of information differ in content, cost, and influence on future decisions.

Displays and Controls. The simulation uses a computer-generated CRT display, illustrated in Figure 5-1. The environment and vehicle are shown as in a moving-map display. Sets of obstacles appear at random positions at the upper edge of the display and move downward at a constant velocity. The operator can move the vehicle symbol horizontally to one of eleven different pathways to avoid the obstacles. He can do so manually, using a joystick, or he can allow the autopilot system to select a path. These control actions are primarily decision making in nature. Dynamics of control are minimized since the obstacle and vehicle velocities are held constant.

The obstacles introduce uncertainty to the task simulation. Each type of obstacle has a specific probability distribution of vehicle loss. Figure 5-2 shows the four obstacle types and their associated probability distributions. For ease of learning, the four obstacle types are designed to be evocative of some obstacles expected to occur in actual control situations -- adverse weather conditions, terrain obstacles, adversaries, and navigational problems.

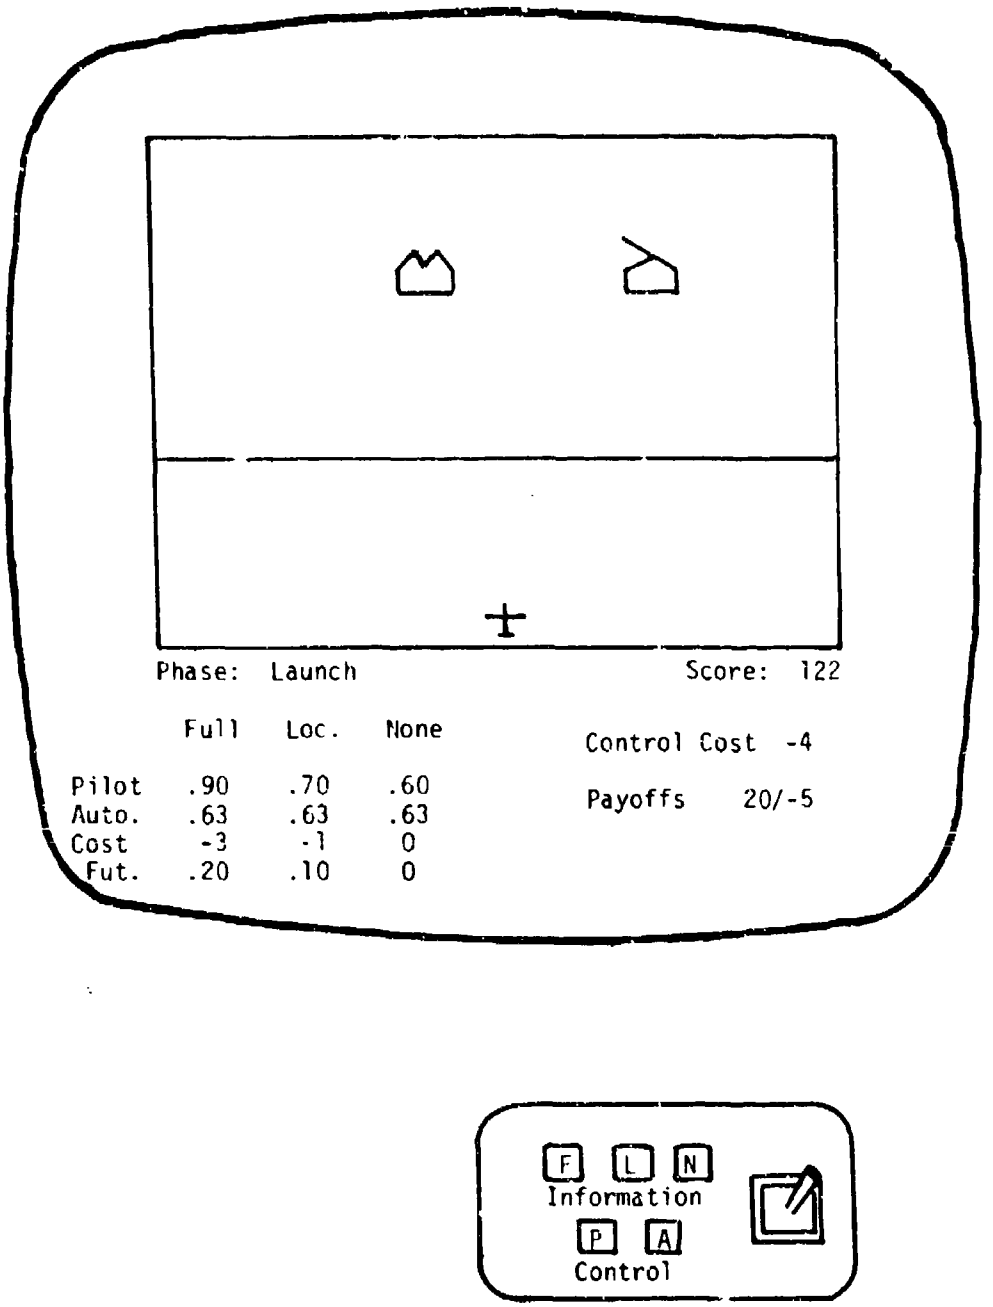

FIGURE 5-1. SIMULATED RPV DISPLAY AND COMMUNICATIONS PANEL

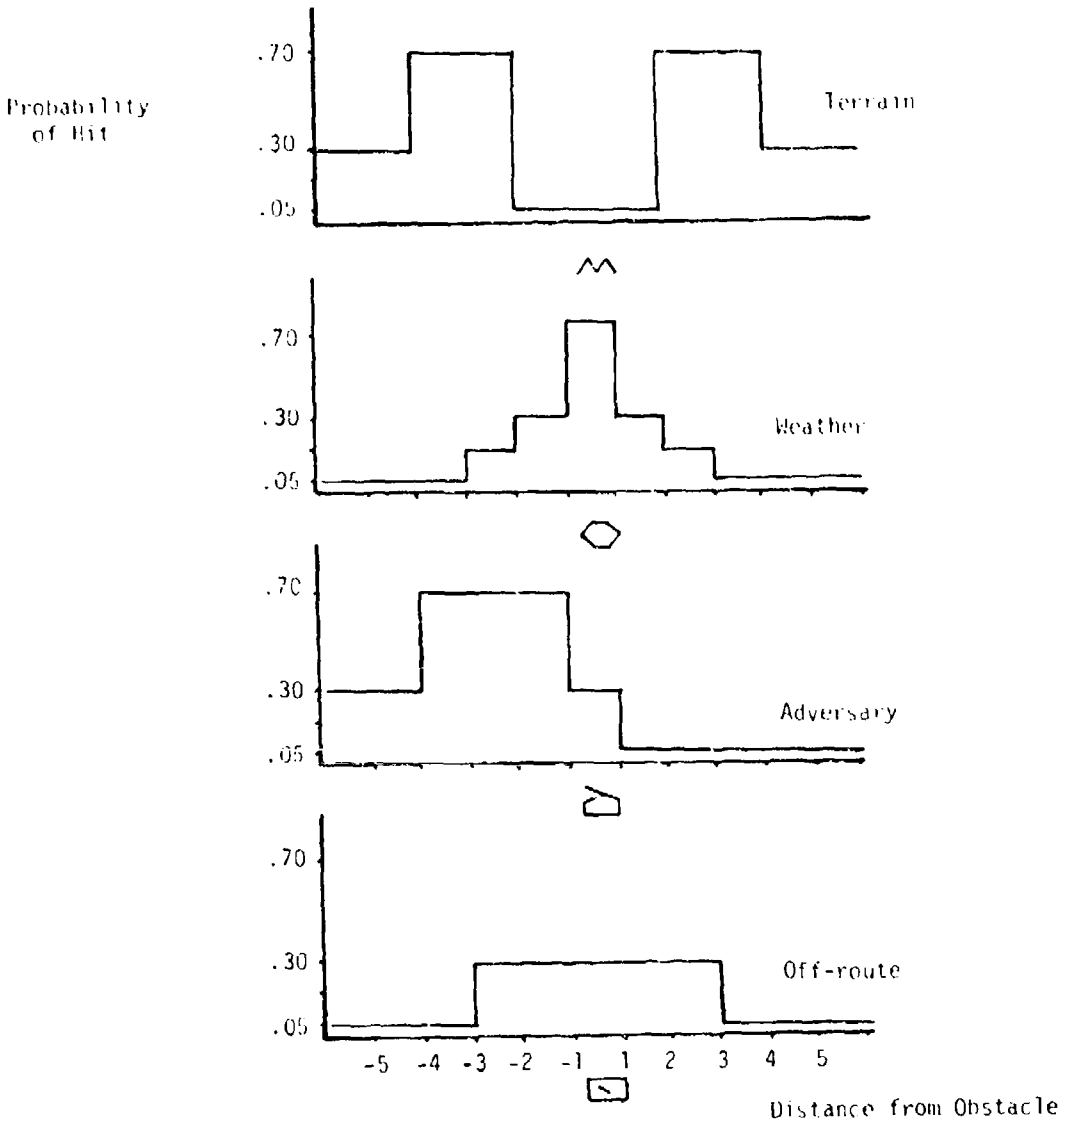

FIGURE 5-2. OBSTACLE HIT PROBABILITY DISTRIBUTIONS

Communications. The identity and location of the obstacles can only be determined from information communicated from the RPV. Three types of information are available for transmission from RPV to operator: (1) Full: A high detail link corresponding to transmission of video images from a nose-mounted camera. This information is simulated by a symbolic presentation of both form and location of the obstacles. While highly informative, this broad band link is costly; (2) Location: A low detail link that provides a report on the location of the obstacles but not their identity. This might correspond to low cost radar information; and (3) None: A minimal information link in which the operator simply relies on flight plan characteristics provided gratis. The presence of obstacles is acknowledged, but data on their identity and location is lacking.

Sequence. The task consists of a series of similar, connected decisions. Prior to the appearance of any obstacles, the operator is appraised of the circumstances surrounding the upcoming decision. He must annunciate (by pressing the button shown in Figure 5.1) the choice of information to acquire. A set of two obstacles is presented at the top of the screen and moves downward. If he selects full information, the differentiated obstacle symbols will move down the screen. If location information is chosen, undifferentiated symbols marking the obstacle locations will pass down the screen. If the minimal information choice is selected, a bar denoting the presence (but not the location) of the set of obstacles moves down. The task moves on continuously, just as an RPV mission does. If the operator does not select an information choice in the time allocated, the minimal information choice is automatically presented.

The operator must make his second decision, that of control allocation, before the obstacle symbols reach the control takeover limit (a line approximately 2/3 of the distance down the screen). This control decision is two-fold: he must decide whether to control the vehicle himself or to give control to the autopilot; and, if controlling himself, he must decide which path to take through the obstacles.

Autopilot Capability. The simulated RPV incorporates an autopilot program capable of limited autonomous control. The autopilot program is subject to error depending on the situational conditions. This combination of autonomous response and unreliability is representative of the behavior associated with both adaptive control systems and preprogrammed response systems. Unpredictable behavior is especially prevalent with adaptive or learning systems since these programs adjust their behavior according to the requirements of the situation. Such systems typically commit errors during changing or unfamiliar conditions.

Unreliability is introduced to the simulated autopilot program through the addition of randomly generated insertions. First, the autopilot selects the optimal action based on the obstacle probability distributions. Then these choices are, with a certain frequency, corrupted by random additions. The frequency of insertion is adjusted according to the task requirements.

Situational Conditions. The various stages of an RPV exercise can be characterized by such factors as difficulty, costs, system reliability, and communications accuracy. Accordingly, the task simulation is configured to involve many of the same factors. It should be noted that these conditions are not experimental variables, but rather are situational conditions designed to provide a demanding exercise for the adaptive model. The experimental variables, defined in a coming section, deal with the form and function of the model. The situational conditions are:

(1) Information Costs. The opening of the communications channel for transmission of information has costs that depend on the level of detail transmitted and on the mission phase. The minimal information choice is gratis, while location and full information have moderate and high costs respectively.

(2) Control Channel Costs. Just as for information transmission, manual control of the RPV is costly. Again, the cost depends on the mission phase.

(3) Information Accuracy. The information transmitted to the operator is subject to error. The unreliability is simulated by random insertions at a given frequency. The frequency of insertion depends on the mission phase.

(4) Future Impact. The establishment of contact with the RPV is assumed to enhance subsequent autopilot capability. Acquisition of full information increases the autopilot reliability on the ensuing decision by a given amount. Acquisition of location information also increases the reliability but does so to a lesser degree.

(5) Payoff Schedule. Payoffs are made independent of whether the operator or autopilot makes the control action. A positive payoff is set for a successful traverse and a negative payoff is associated with an unsuccessful one. Each of these are constants dependent on the mission phase.

The presentation of conditions is organized into three distinct mission phases -- launch, enroute, and target. These phases are similar to the types of exercises an RPV must perform, thus lending a degree of surface credibility to the task. Each phase has set levels of payoffs, information accuracy, autopilot capability, and future impact. For diversity, the communications costs of obtaining information and control are randomly varied about central values established for each phase. The mission phases and their associated conditions are:

(1) Launch--Low risk flight in safe area.

Low: Communications costs; payoffs

High: Information accuracy, autopilot capability, future impact

(2) Enroute--Traverse of hostile region.

Low: Information accuracy, autopilot capability

High: Communications costs, payoffs, future impact

(3) Target--Critical approach to target.

Low: Information accuracy, communication costs, autopilot capability, future impact

High: Payoffs

The conditions are displayed to the operator prior to each decision. The conditions shown in Figure 5-1 are an example of a launch phase decision. The upper line on the left provides estimates of the operator's probability of success under each form of information. The next line indicates the probability of success of the autopilot. This estimate does not change with the information acquired. The next two lines are the information costs and the future impact. The information costs are the points subtracted for acquisition of each type of information. The future impact is the degree of augmentation of the autopilot on the succeeding decision. The right side of the display shows the control cost for pilot control along with the payoffs for successful and unsuccessful actions.

5.4 Decision Model

The experimental situation results in a simplification of the general information and control model of Section 4-2. The decision space is reduced to three information choices and two subsequent control choices.

Nevertheless, a variety of multidimensional consequences are possible, stemming from the various combinations of costs, payoffs, and actions possible.

The decision tree shown in Figure 5-3 summarizes the communications choices. The initial portion of the more general tree depicted earlier in Figure 3-3, the information gathering decision by the autopilot for its own use, is not represented in this tree. The experimental situation corresponds to the case where the autopilot has continuous, cost-free access to high detail information. This allows the vehicle-to-operator communications to be emphasized in this initial study.

The major decision faced by the operator is thus the choice of information and control to transmit. All six combinations of information and control are shown in the figure. Two combinations of special note are those of full information/autopilot control and location information/ autopilot control. These combinations might be selected when the operator wishes to maintain supervision of the remote system activities (and increase autopilot reliability) without taking control. The remaining combinations have more obvious rationales.

The probability of occurrence of each of the 12 outcome catagories can be estimated analytically or from performance histories. The probability of success for the autopilot can be derived from the obstacle probability distributions, the autopilot unreliability level, and the future impact. This calculation is the following:

$$P(s) = \{1-R-F)\} \cdot \max_{a} \{(1-P(H_b|a))(1-P(H_c|a))\}$$
$$-(R-F) \cdot \frac{1}{11} \sum_{a=1}^{11} (1-P(H_b|a))(1-P(H_c|a)) \qquad (5\text{-}1)$$

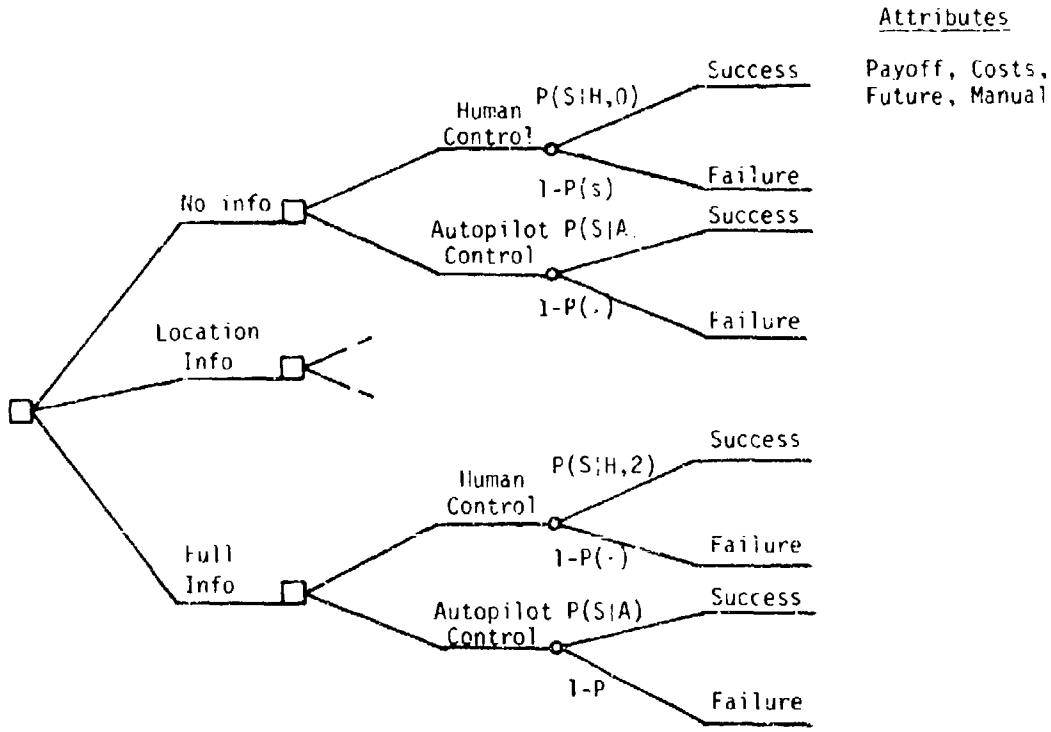

FIGURE 5-3. STRUCTURE OF COMMUNICATIONS DECISIONS

where

a is an action; the choice of one of the eleven paths

R is the reliability of the autopilot; the percent of random actions

F is the influence of preceeding communications on increasing autopilot reliability

$P(H_b|a)$ is the probability of a hit from obstacle b given that action a is taken (see Figure 5-2 for hit distributions)

$P(H_c|a)$ is the probability of a hit from obstacle c given action a

The first of the two factors in Equation 5-1 is simply the probability of success of traversing the two obstacles encountered, given that the highest probability path is taken. The second factor takes into account the autopilot reliability. R-F represents the frequency of random actions by the autopilot, and the summation provides an averaging of the effectiveness of these responses across all possible actions. The probability of success so determined for the autopilot is the same for all information choices.

The probability of success of operator control, on the other hand, depends strongly on the type of information displayed. A different estimator must be used for each of the three levels of information transmission. The specific actions and states are not analyzed, since the operator may not act optimally. The operator is simply assumed to have different distributions of outcomes under the different information sources and mission phases. Separate estimators, based on moving averages of the success frequencies, are maintained for each of the 9 combinations of information type and mission phase.

Decision Attributes. A preliminary set of attributes has been selected that appears representative of the types of factors that enter into RPV communications decisions. An attempt was made to choose attributes by the criteria described in Section 3.3--accessibility, monotonicity, completeness, and independence. The factors and their derivation follow:

(1) Communications Costs. The direct costs associated with use of the communication channel--energy expenditures, detection, etc. In the simulation, the combined costs of information transmission and manual control comprise this scale.

(2) Control Outcomes. The actual consequences of control. In the real world, these consequences may include loss of vehicle, system damage, fuel depletion, political gains, or attainment of goal. The consequences here are defined simply as success and failure. The expected payoff associated with success and failure determine the attribute level.

(3) Future Impact. The effect on future decisions resulting from the establishment of communications. This factor is representative of the many indirect effects of RPV communications, among them such factors as control continuity and influence on autopilot effectiveness. Future impact is defined in this exercise as the augmentation of autopilot capability on the succeeding decision. As was described earlier, greater amounts of man-machine interaction lead to reductions in autopilot unreliability on the immediately succeeding decision. The attribute level is defined as the percent improvement resulting from the communication.

(4) Control Preference. The subjective preference for manual over autopilot control. This factor is one of a number of possible purely subjective factors--manual control propensity, operator loading, concrete versus abstract preference, etc. The attribute level in this study is binary, as it corresponds to the presence or absence of manual control in the alternative.

The attributes thus represent four main sources of consideration in communications decisions--costs, direct consequences, indirect affects, and operational preferences. The same attributes are used to describe both the information seeking and the control allocation decisions. The evaluation of each of the 6 combinations of information and control would be made according to the following equation:

$$MAU_j = \sum_{i=1}^{4} W_i \cdot x_{ij}$$

where

MAU_j is the aggregate (multi-attribute) utility of alternative j
x_{ij} is the level of attribute i in alternative j
W_i is the inferred weight of attribute i

The information/control choice with the highest MAU would be selected by the model. If the operator selected a different alternative, the model would be adjusted according to the methods described in Section 3.4. The absolute rule is used for adjustment.

5.5 Experimental Procedure

5.5.1 Experimental Variables. The experimental hypotheses deal with the effectiveness of prediction and degree of aiding provided by the model. Accordingly, the following experimental variables and levels are planned:

(1) Model Form--Two Levels

Differential weighting. Use of model inferred attribute weights for prediction, analysis, and aiding.

Unity Weighting. Control condition in which arbitrary (all 1.0) weights are used for prediction and aiding. All of the weights are defined to be positive except cost.

(2) Aiding--Two Levels

Model Based Aiding. Operator makes information and control choices after observing model recommendations.

No Aiding. Operator makes information and control choices without model recommendation.

The four combinations of the above experimental conditions provide an essentially complete testing sequence for the experimental hypotheses. The model form conditions provide a basis for testing the predictability and validity of the adaptive model, while the complete set of conditions allow testing of the influence of aiding on performance. The aiding conditions also result in an indication of the degree of operator acceptance of the machine recommendations.

5.5.2 Performance Measures. The standard performance measures such as errors and speed only partly describe the quality of performance in a shared control task. The close coupling of man and semi-autonomous machine require additional evaluations of individual contributions, decision model performance, and decision quality.

System Performance. The overall system performance is described using a single index, the score. The score is derived from the number and cost of errors committed and the communications costs expended:

SCORE = {PAYOFFS} - {PENALTIES + COMMUNICATION COSTS}

The communications costs include both the costs of information transfer and the costs of assumption of manual control. The score is presented to the subject as a single index of performance, and his compensation depends to a large extent on the measure. The complexities of having speed as a second, independent measure are avoided by presenting the task at a set pace.

Decision Model Performance. The effectiveness of the decision model in inferring decision parameters and predicting operator behavior can be determined by a number of methods. Among these means of model validation are axiomatic tests, measures of prediction, construct validity tests, and checks of operator acceptance. Prediction is the simplest of these. The ability of the adaptive model (and of a utility weight model) to predict behavior in both the information and control decisions can be determined directly. Construct validity tests are more difficult. These tests are made by comparing the inferred weights with weights estimated off-line by other techniques. The off-line estimation techniques may involve direct estimation, paired-comparisons, or the Ynetma-Torgerson "interpolation between the corners" method. The last of these, the Yntema-Torgerson technique, was selected for use in the study because of its simplicity and reliability.

Auxiliary Measures. Additional measures of decision quality and decision consistency were implemented. The decision quality measures are determinations of the deviation from maximum expected utility exhibited

by the operator. Assuming model accuracy, this is a measure of sub-optimality of behavior due to logical inconsistency. The second measure, decision consistency, is measured by the overall stability of the estimated preference structure.

5.5.3 Experimental Design and Procedures. The hypotheses were examined using a three factor experiment with repeated measures. The three factors were the model, level of aiding, and sequencing of conditions. Repeated measures were taken across the first two factors as shown in Figure 5-4. In this design, each subject was exposed to all four combinations of conditions.

The eight subjects participating in the study were recruited from nearby Air National Guard units. They represented the type of personnel who might interface with computer-aided communications system. The subjects' ages ranged from 22 to 42. Four of the subjects were pilots and three had extensive experience with computer systems. All had some college experience (1 to 12 years). The eight subjects were assigned randomly to the four groups.

Each subject performed the task during three sessions of 2 hours duration. The first hour of the first session served as a familiarization and practice period. The subjects were given instructions on system operation and were provided hands-on experience with the equipment. The remaining hour of the first session, both hours of the second session, and the first hour of the final session were devoted to experimental runs. Each of the experimental runs lasted 55 minutes and consisted of four complete sequences of launch, enroute, and terminal phases. The subjects were paid on an hourly basis and received up to $4.00 an hour additional as a bonus contingent on performance. The bonus was a function of their score compared to that of the other participants.

	No Aiding		Aiding	
	Unity Weights	Infrared Weights	Unity Weights	Infrared Weights
Group 1	C	D	A	B
Group 2	D	C	B	A
Group 3	B	A	C	D
Group 4	A	B	D	C

(Letters denote order)

FIGURE 5-4. EXPERIMENTAL DESIGN

The subjects were informed that on two of the runs, aiding would be provided in the form of communications recommendations. The subjects were not informed as to the nature of the aiding. The experimental runs ended with an off-line estimation of each subjects' attribute weights using the Yntema-Torgerson technique.

Data Report. Following every experimental segment, the computer provided a printout of the experimental performance indices. Segments consisted of ten decisions, spanning a single mission phase. The 24 seconds required for the printing permitted a short rest period. The data report provides a breakdown of the performance score into that attributable to the operator and to the autopilot. Also, analyses of the information seeking behavior and model performance are included. A typical data report is shown in Figure 5-5.

```
MEAN ATTRIBUTE LEVELS   5 550E 0 3 775E 0 2 625E 0 4 781E 0 4 256E 0 3 731E 0
MEAN WEIGHT  -1 000E 0 1 000E 0 1 000E 0 1 000E 0
NO  OF INFO  & CONT  REQUESTS:      2     0     1     0     2     5
-----LAUNCH PHASE------
NO  OF OPERATOR CONTROLS        3
NO  OF MACHINE CONTROLS:       7
NO  OF OPERATOR ERRORS:        0
NO  OF MACHINE ERRORS:       0
NO  OF INFO  OVERIDES:       1
NO  OF CONTROL OVERIDES:       2
PERCENT OPERATOR ERRORS:       0
PERCENT MACHINE ERRORS:        0
MEAN OF PROBABILITIES   8 000E-1     0E 0 5 000E-1     0E 0 6 000E-1 5 000E-1
MEAN MAXIMUM EV   1 410E 1
MEAN DEVIATION OF MAXIMUM EU    3 900E 0
EXPECTED VALUE FOR DECISION:  1 493E 1
COMM  COST UNDER OPERATOR CONTROL:        14
COMM  COST UNDER MACHINE CONTROL:         4
TOTAL ERRORS:        0
TOTAL COMM  COST:       18
SCORE     182
DECISION PERCENT PREDICTED:       80
MEAN INFORMATION VALUE:   2 750E 0 1 250E 1-3 200E 1-1 250E 1
```

FIGURE 5-5. SAMPLE DATA REPORT

6. EXPERIMENTAL RESULTS AND DISCUSSION

6.1 General Observations

The choices presented in the RPV task simulation were found to be sufficiently varied and difficult to provide a good initial exercise for the decision model. A wide variety of behaviors were observed and modeled. The task simulation was also found to be sufficiently demanding to maintain a high level of subject interest. The subjects learned the task procedure readily and by the end of the training session, could efficiently handle the task requirements.

6.2 Model Descriptive Performance

6.2.1 Prediction. Table 6-1 shows the percent of decisions predicted by the experimental (inferred weight) model and by the control (unity weight) model. Each experimental session resulted in a data point of percent prediction. The prediction itself was the choice of one of the six possible combinations of information and control.

The experimental model was found to predict behavior significantly more effectively than the unity weight model ($F=10.1$, $df=1,4$, $p<.05$). Overall, a 50 percent prediction rate was observed with the adaptive model, as compared with a 40 percent prediction rate by the unity model. The seemingly low rates of prediction were due to the difficulty of the information and control choices. Dominated alternatives and clear choices were seldom present. Aiding by recommendation using either model, also improved the prediction accuracy ($F=26.2$, $df=1,4$, $p<.01$). Prediction by the unity weight model increased from a level of approximately 35 percent to one of 45 percent, while for the inferred weight model, aiding increased the prediction rate from a 45 percent level to a 56 percent level. These increases were presumably due to acceptance of the aiding by the subjects.

TABLE 6-1. MEANS OF PREDICTION MEASURES FOR THE COMBINATIONS OF MODEL AND AIDING

	Non-Aiding		Aiding	
	Control Model	Experimental Model	Control Model	Experimental Model
Percent of Communications Decisions Predicted	34.6	44.6	44.9	55.7
Percent of Information Decisions Predicted	58.3	65.3	70.5	76.4
Percent of Control Decisions Predicted	73.2	71.3	76.1	82.9

The majority of the prediction increment between the inferred and unity weight models was traceable to increased accuracy of prediction in the information choice. As shown in Figure 6-1, a significant increase in information choice prediction was found with the inferred model as compared to the unity model ($F=17.2$, $df=1,4$, $p<.01$). Also, increases in information prediction were found with the aiding conditions compared to the non-aiding conditions ($F=17.6$, $df=1,4$, $p<.01$). Neither the modeling or the aiding differences reached significance with the subsequent control choice. However, an interesting interaction effect was found with the control choice behavior, as shown in Figure 6-1. Acceptance of model recommendations, indicated by an increase of prediction with aiding, was present only with the experimental model ($F=7.7$, $df=1,4$, $p<.05$). The unity weight model did not seem to engender such acceptance.

6.2.2 Decomposition. The components of the multi-attribute model were also found to be predictive of components of behavior. The costs actually expended by the subjects were highly correlated with the communications cost weight ($r=.853$, $p<.001$) and moderately correlated with the control outcome weight ($r=-.51$, $p<.05$). These findings indicate an intercorrelation among attributes, pointing to a need for more careful initial definition of the attributes. The intercorrelation appears to be due to the fact that costlier information is typically associated with higher probabilities of success. Nevertheless, the robustness of the linear model allows it to function well regardless. The number of requests for information having non-zero future impact was similarly correlated with both the future impact weight ($r=.51$, $p<.05$) and the communications cost weight ($r=.57$, $p<.05$). Here the acquisition of information having high future impact was associated with increased costs.

The relationship of behavior to model-inferred weights can be illustrated by example. Figure 6-2 shows three representative samples of divergent behavior and the corresponding weight vectors. Each example

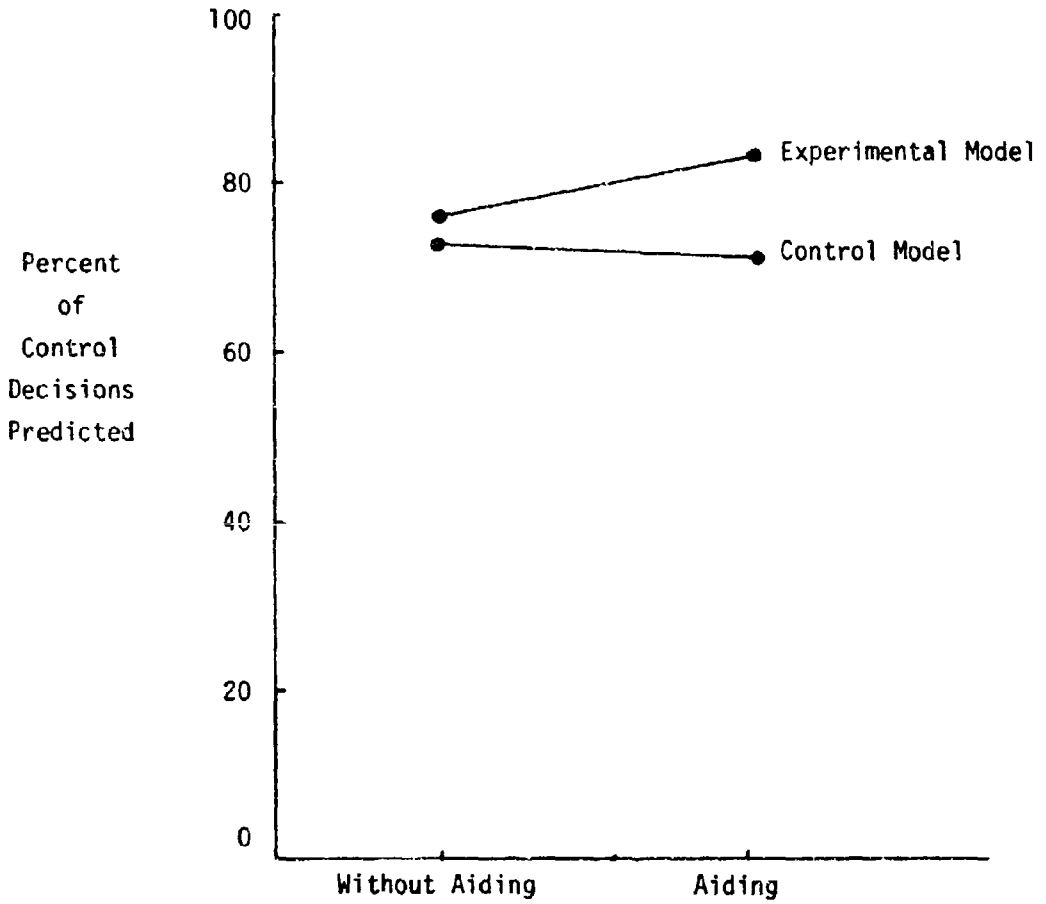

FIGURE 6-1. INTERACTION OF AIDING WITH MODEL FORM

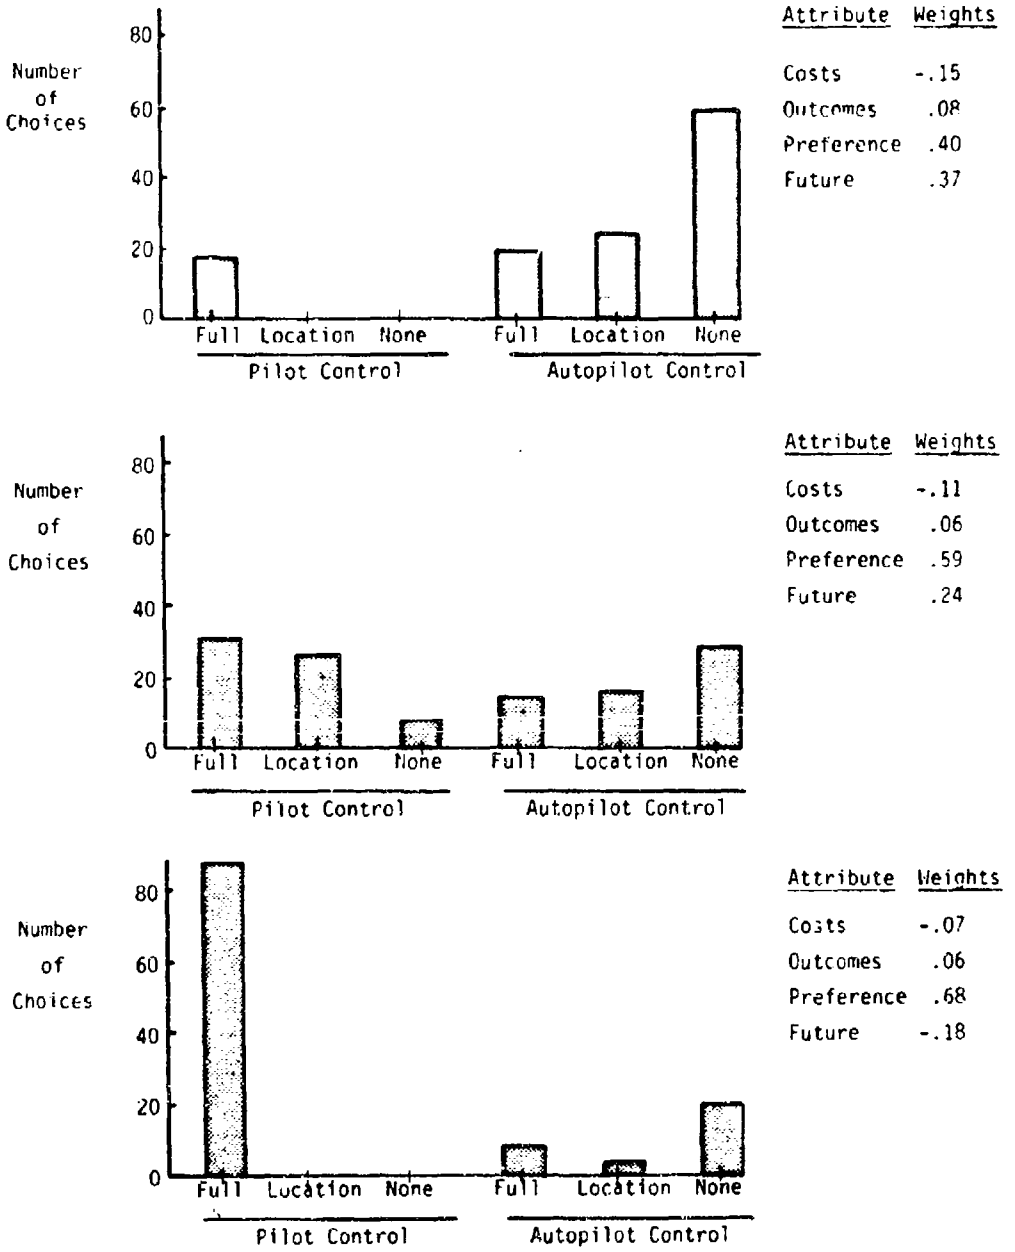

FIGURE 6-2. REPRESENTATIVE SAMPLES OF CHOICE DISTRIBUTIONS AND THE ASSOCIATED MODEL WEIGHTS

depicts the distribution of choices made during an experimental session. The averaged weight estimations for each session are shown to the right of the figures. The upper figure is from a session during which the subject minimized costs by emphasizing autopilot control. This behavior is reflected in the relatively high negative weight for communications costs and the low weight for manual control preference. The moderate level of the future impact weight appears to be due to the frequent acquisition of full and location information during autopilot control. It should be noted that while the weight vectors are normalized (the absolute levels sum to one), the relative contributions of the attributes are not proportional to the weights, since the ranges of the attributes are different.

The second histogram of Figure 6-2 shows a more balanced behavior. The information and control choices are distributed more equally among the options available. The communications cost weight is less extreme than in the first case, and the manual control preference is somewhat higher.

An example of strong manual control propensity is shown in the lower figure. Here the communications cost weight is quite low, as might be expected, with frequent acquisition of full information and manual control. Also, the manual control preference is high and the future impact weight is negative. The negative weight for future impact may be due to the subject's neglect of location information in both manual and autopilot modes.

6.3 Task Performance

No significant differences in the performance score were noted between the aided and unaided conditions with either model. This is understandable, since the adaptive model does not direct the operator to maximize score, but rather attempts to capture, analyze, and extrapolate his behavior to new situations.

Nevertheless, the adaptive model was useful in identifying behaviors which led to superior performance. Corrections between scores achieved and the inferred attribute weights indicated that the communications costs and control outcome attributes were the factors of primary importance. The control preference and future impact attributes appeared to be more subjective and less consequential regarding performance. The score attained correlated very highly with the communications cost weight ($r=-.853$, $p<.01$) and moderately so with the control outcome weight ($r=.448$, $p<.05$). Between these two factors, roughly 93 percent of the variance in scores was accounted for. Inspection of scatter plots of the individual attribute weights with the scores attained, indicated that inverted-u relations were present between each of the four weights and the score. This seems reasonable, as an optimal region for each importance weight would be expected. The sizable linear relationship of score attained to the communications cost and control outcome weights appeared be due to the concentration of scores on the low end of these scales. subjects seemed to take sufficient account of these factors. Instead, appeared to over-emphasize the control preference and future impact factors.

An interesting finding was the relationship of score achieved to consistency of behavior with respect to the model. With the unity weight model, a moderate correlation was observed between the deviation from expected utility (DEU) and the score ($r=.52$, $p<.05$). The DEU is essentially a distance measure reflecting closeness of behavior to the model recommendations. Typically, the lower the DEU, the higher was the observed score. This relationship of decision consistency to score was also seen with the adaptive model during aiding ($r=.80$, $p<.05$). It appears that both the unity and the adaptive models resulted in improvements in performance when deviations from the model recommendations were minimized.

6.4 Validation

A preliminary check on the model validity was made by comparing the inferred parameters with weights estimated through off-line procedures. The Yntema-Torgerson "Interpolation Between the Corners" technique was employed as the comparison method. In this, each subject estimated on a scale of zero to 100 the attractiveness of various hypothetical information and control choices. The scale was anchored at the zero and 100 points by the worst and best combinations of conditions, respectively. These combinations specify 2 of the 16 possible combinations of the 4 attribute extremes. The remaining 14 combinations of corner conditions were presented in sequence to the individual subjects. The resulting ratings were then normalized so as to be comparable in scale to the model inferred weights (see Sheridan and Ferrell (1974) for a description of the derivation procedures).

A very strong test of the similarity of the eight pairs of attribute profiles was made using a two factor ANOVA with repeated measures on both factors (estimation methods and attributes). The test that the two profiles were identical was rejected ($F=5.43$, $df=3,21$, $p<.01$). Nevertheless, correlation coefficients between the attribute estimations by the two methods averaged .46, which is significant at the [illegible] level.

Comparisons of the on-line and off-line methods of estimation were also performed by correlation with behavior. As noted previously, the adaptively inferred weights for communications costs and future impact correlated significantly with the costs expended and with the frequency of information acquisition in the task. The off-line estimates of the future impact weights also correlated significantly with information acquisition ($r=.56$, $p<.05$), but the critical weight for communications

cost did not correlate with costs expended. Also, correlations between the communication cost weight and the task performance did not reach significance. From these findings, it appears that the adaptive estimation procedures had an advantage in prediction over the off-line technique.

7. CONCLUSIONS AND RECOMMENDATIONS

7.1 Adaptive Decision Modeling

The present study demonstrates some of the potential of on-line, adaptive techniques for modeling communications decisions. The information seeking decisions involved in remotely piloted vehicle supervision were seen to be amenable to analysis using multi-attribute decision models. In the RPV situation, the operator is repeatedly required to make complex, subjective decisions regarding information and control options. Multi-attribute models using pattern recognition techniques for estimation were seen to be able to capture much of this behavior.

The preliminary experimental studies demonstrated the speed, simplicity, and robustness of the adaptive technique. The on-line estimation technique was found to be more predictive of behavior than either the unity weight or off-line methods tested. The adaptive model was also useful in identifying differing decision policies and partitioning out components of behavior, at least to a rudimentary level. Finally, the adaptive model appeared to be accepted to some degree by the operators, since the model prediction rate increased with display of the model's recommendations.

Of course, the multi-attribute models and adaptive estimation procedures are not proferred as the general methodology for communications decision modeling. These techniques are specific to decisions that are complex, subjective, and recurrent. Some rough criteria are given identifying situations that favor use of off-line parameter estimations, unity weight models, linear cue formulations, and related techniques. The coming studies will attempt to refine these guidelines.

7.2 Decision Modeling and Task Performance

The adaptive model was found to be helpful in identifying strategies which led to superior performance. The time-proven capabilities of linear models in analyzing components of performance were again seen.

Aiding, based on recommendations by the on-line model, did not produce the significant increases in performance expected, although those subjects who followed the recommendations most closely achieved the highest scores. The effect on these subjects appeared to be the classic reduction of randomness or noise in behavior. Strong improvement in performance with aiding are expected to be more likely in situations of greater time stress and decision complexity.

7.3 Value of Information

The probabilistic multi-attribute model provides an ideal framework for ascertaining the value of information. The benefit of an information system in a set of task situations can be determined by aggregating the constituent influences of the communications.

The adaptive model contributes to this analysis by providing estimates of the model parameters in operational situations. A figure of merit can then be given to a specific information system by aggregating the values derived over the distribution of situations. Such a procedure is planned for the coming study, along with sensitivity analyses of the various model parameters and situational factors. The resulting methodology should be useful for evaluating alternative systems of information sensing, processing, encoding, transmitting, and display. Also, the techniques should be helpful for specifying information needs and training operators to make effective communications decisions.

7.4 Management of Communications

The availability of a methodology for information value determination opens up the possibility of managing communications from an automated system. Ideally, the operator would then be appraised of only essential information, instead of having to frequently and inefficiently interrogate the system.

The acceptance of model recommendations noted in this study is encouraging regarding the possibility of remote management of communications. The next step, scheduled for the coming year, is the model-directed presentation of communications to the operator.

8. REFERENCES

Anderson, N.H. and Shanteau, J.C. Information Integration in Risky Decision Making. Journal of Experimental Psychology, 1970, 84:441-51.

Baron, S. Some Comments on Parameter Identification in the Optimal Control Modes. Systems, Man and Cybernetic Review, February 1977, 6(1):4-6.

Bauerschmidt, D.K. and La Porte, H.R. Techniques for Display/Control System Integration in Supervisory and Monitoring Systems. In Sheridan, T.B. and Johannsen, G. Monitoring Behavior and Supervisory Control. New York: Plenum Press, 1976.

Beach, B.H. Direct and Indirect Methods for Measuring Utility. Scientific Interim Report, Department of Psychology, University of Washington, ONR N00014-67-A-0103-0011/NR151-313, Personnel and Training Research Programs Office, Office of Naval Research, July 4, 1972.

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

Becker, G.M., DeGroot, M.H., and Marschak, J. Stochastic Models of Choice Behavior. Behavioral Science, 1963, 8:41-55.

Bejczy, K. Machine Intelligence for Autonomous Manipulation. In Heer, E. (Ed.) Remotely Manned Systems. California Institute of Technology, Los Angeles, CA: No. Hollywood Printing Co., 1973.

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

Brunswik, E. The Conceptual Framework of Psychology. Chicago: University of Chicago Press, 1952.

Brunswik, E. Thing Constancy as Measured by Correlation Coefficients. Psychological Review, 1940, 47:69-78.

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

Duda, R.O. and Hart, P.E. Pattern Classification and Scene Analysis, Wiley-Interscience: New York, 1973.

Edwards, W. and Guttentag, M. Experiments and Evaluations: A Re-Examination. In Bennett, C. and Lumsdaine, A. (Eds.) Experiments and Evaluations. New York: Academic Press, 1975.

Einhorn, H. and Hogarth, R. Unit Weighting Schemes for Decision Making. Organizational Behavior and Human Performance, 1975, 13:171-192.

Emery, J.C. Organizational Planning and Control Systems. New York: MacMillan Company, 1969.

Engstrom, K.D. and Rouse, W.B. Telling a Computer How a Human Has Allocated His Attention Between Control and Monitoring Tasks. Proceedings of the Annual Conference of Manual Control, 1977.

Pelson, J. Cybernetic Decision Systems. Louisiana: CDS Publishing Co., 1975b.

relson, 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.

Ferrell, W.R. Command Language for Supervisory Control of Remote Manipulation. In Heer, E. (Ed.), Remote Manned Systems, California Institute of Technology, Los Angeles, CA: No. Hollywood Printing Co., 1973.

Fischer, G.W. Four Methods for Assessing Multi-Attribute Utilities: An Empirical Validation. Engineering Psychology Laboratory, University of Michigan (Ann Arbor, MI), Technical Report No. 037230-6-T, ONR N00014-67-A-0181/NR197-014, Code 455, Office of Naval Research, September 1972.

Fisher, W. Remote Viewing System. In Heer, E. (Ed.) Remotely Manned Systems, California Institute of Technology, Los Angeles, CA: No. Hollywood Printing Co., 1973.

Fogel, L.S., Englund, C.E., Mout, M.L. and Hertz, T.P. Principles of Display and Control Design for Strike RPV's. Decision Science (San Diego, CA) Technical Report, Control N00014-72-C-0196, June 1974.

Freedy, A., Davis, K.B., Steeb, R., Samet, M.G. and Gardiner, P.C. Adaptive Computer Aiding in Dynamic Decision Processes: Methodology, Evaluation, and Application. Perceptronics, Inc. (Woodland Hills, CA) Final Technical Report PFTR-1016-76-8/30, Contract N00014-73-C-0286/NR 196-128, ARPA Order No. 2347, Code 455, Office of Naval Research, August 1976.

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., Weisbrod, R.L., and Steeb, R. A Pattern Recognition Approach to Subjective Value Estimation for Task Control Allocation. Journal of American Society for Cybernetics, 1975.

Freedy, A., Weisbrod, R., and Weltman, G. Self-Oprimization 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.

Fu, K.S. Sequential Methods in Pattern Recognition and Machine Learning. New York: Academic Press, 1970.

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

Gardiner, P.C. & Edwards, W. Public Values: Multi-Attribute Utility Measurements for Social Decision Making. In Schwartz & Kaplan (Eds.) Human Judgment & Decision Processes: Formal and Mathematical Approaches. New York: Academic Press, Inc., 1975.

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.

Goldberg, L.R. Simple Models or Simple Processes? American Psychologist, 1968, 23:483-495.

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

Goodman, C. Direct Estimation Procedures for Eliciting Judgments about Uncertain Events. Engineering Psychology Laboratory, University of Michigan (Ann Arbor, MI) Technical Report 011313-5-T, 1973.

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

Huber, G.P. Methods for Quantifying Subjective Probabilities and Multi-Attribute Utilities. Decision Sciences, 1974, 5:430-458.

Johansson, G. Preview of Man Vehicle Control Session. In Sheridan, T.B. and Johansson, G. (Eds.) Monitoring Behavior and Supervisory Control. New York: Plenum Press, 1976.

Keeney, R.L. and Sicherman, A. An interactive Program for Assessing and Analyzing Preferences Concerning Multiple Objectives. International Institute for Applied Systems Analysis Report RM-75-12, Laxenburg, Austria. April 1975.

Keeney, R.L., and Raiffa, H. Decision Analysis with Multiplte Conflicting Objectives, Preferences, and Value Trade-Offs. Wiley: New York, 1975.

Kneppreth, N.P., Gustafson, D.H., Johnson, E.M., and Leifer, P. Techniques for the Assessment of Worth. U.S. Army Research Institute for the Behavioral and Social Sciences, Technical Paper 254, Contract DAHC19-72-0026, December 1973.

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

Luce, R.D. and Raiffa, H. Games and Decisions. New York: Wiley & Sons, 1957.

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, Vol. III. New York: Wiley, 1965.

MacCrimmon, K.R. and Taylor, R.N. Decision Making and Problem Solving Processes. In Dunnett, M.D. (Ed.) Handbook of Industrial and Organizational Psychology. New York: Rand McNally, 1972.

Marschak, J. The Payoff-Relevant Description of States and Acts. Econometric, October 1973, 21(4):719-725.

McKendry, J.M., Enderwick, T.P., and Harrison, P.C. A Subjective Value Approach to Information Utility. Human Factors, 1971, 13(6):503-509.

Miller, A.C., Merkhofer, M.W., Howard, R.A., Matheson, J.E., and Rice, T.R. Development of Automated Aids for Decision Analysis. Stanford Research Institute Technical Report, May 1976.

Miller, George A. The Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information. Psychological Review, 1956, 63, 81-07.

Mills, R.G., Bachert, R.F., and Aume, N.M. Summary Report of AMRL Remotely Piloted Vehicle (RPV) System Simulation Study II Results. Ohio, AMRL Technical Report AMRL-TR-75-13, Wright-Patterson Air Force Base, February 1975.

von Neumann, J. and Morgenstern, O. Theory of Games and Economic Behavior. Princeton: Princeton University Press, 1947.

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.

Raiffa, H. Preferences for Multiattributed Alternatives. Rand Corporation (Santa Monica, CA) Memorandum RM-5868-DOT/RC, April 1969.

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

Rouse, W.B. Human Interaction with an Intelligent Computer in Multi-Task Situations. Proceedings Annual Conference on Manual Control, NASA Ames Research Center, May 1975.

Roscoe, S.N. and Eisele, J.E. Integrated Computer-Generated Cockpit Displays. In Sheridan, T.B. and Johannsen, G. (Eds.) Monitoring Behavior and Supervisory Control. New York, Plenum Press, 1976.

Samet, M.G. Utility Assessment by "Interpolation Between Corners: An Application to the Evaluation of Drug Efficacy." U.S. Army Research Institute for the Behavioral and Social Sciences (Arlington, VA) Unpublished Manuscript, December 1976.

Savage, L.J. The Foundations of Statistics. New York: Wiley & Sons, 1954.

Shannon, C.E. and Weaver, W. The Mathematical Theory of Communication. University of Illinois Press, Urbana, Illinois, 1949.

Sheridan, B. Toward a Model of Supervisory Control. Sheridan, T.B. and Johannsen, G. (Eds.), Monitoring Behavior and Supervisory Control, New York: Plenum Press, 1976.

Simon, H.A. Models of Man. New York: Wiley & Sons, 1957.

Singleton, W.T. The Model Supervisor Dilemma, Sheridan, T.B. and Johannsen, G. (Eds.) Monitoring Behavior and Supervisory Control. New York: Plenum Press, 1976.

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

Slovic, P., Fischhoff, B., and Lichtenstein, S. Behavioral Decision Theory. Annual Review of Psychology, 1977, 28:1-39.

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.

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.

Steeb, R. and Freedy, A. Man-Machine Interaction in Adaptive Remote Systems, IEEE Conference on Cybernetics and Society, Washington, D.C. November 1-3, 1976.

Tversky, Amos and Kahneman, Daniel. Judgment Under Uncertainty: Heuristics and Biases. Office of Naval Research, Report NORIB 13(1), Oregon Research Institute ONR N00014-73-C-0438/NR 197-026, ARPA Order #2449, August 1973.

Tversky, A., Slovic, P. and Lichtenstein, S. Subjective Optimality: Conceptual Issues, Empirical Results, and Decision Aids. Office of Naval Research Conference on Applications of Decision and Information Processing Research, Monterey, California, October 4, 1972.

Wendt, D. Value of Information for Decisions. Journal of Mathematical Psychology, 1969, 6:430-443.

Wendt, D. Some Criticisms of the General Models used in Decision Making Experiments. Engineering Psychology Laboratory, University of Michigan (Ann Arbor, Michigan) Report No. 011313-8-T, ONR, N00014-67-A-01810049/NR 197-21, ARPA Order #2105, Office of Naval Research, 2 November 1973.

Whittemore, J. and Yovits, M.C. A Generalized Conceptual Flow for the Analysis and Flow of Information. Journal of the American Society for Information Science, May-June 1973.

Whittemore, J. and Yovits, M.C. The Quantification of Information Used in Decision Processes. Information Sciences, 1974, 7:171-184.

von Winterfeldt, D. An Overview, Integration, and Evaluation of Utility Theory for Decision Analysis. Social Science Research Institute, University of Southern California (Los Angeles, CA), Research Report 75-9, August 1975.

von Winterfeldt, D. and Fischer, G.W. Multi-Attribute Utility Theory: Models and Assessment Procedures. Engineering Psychology Laboratory Technical Report 011313-7-T, November 1973.

Yntema, D.B., and Torgerson, W.S. Man-Computer Cooperation in Decisions Requiring Common Sense. IRE Transactions on Human Factors in Electronics, 1961, HFE-2:20-26.

Zimmerman, R. Graphic Video Displays for Process and Man-Vehicle Control. In Sheridan, T.B. and Johannsen, G. (Eds.) Monitoring Behavior and Supervisory Control. New York: Plenum Press, 1976.

DISTRIBUTION LIST

AMRL/HE (Lt. Col. J. Birt)
Wright-Patterson AFB OH. 45433

AMRL/HE (Dr. Topmiller)
Wright-Patterson AFB OH. 45433

AMRL/HE (Dr. Mills)
Wright-Patterson AFB OH. 45433

Dr. Martin Tolcott
Office of Naval Research
800 N. Quincy St.
Arlington, VA. 22217

Dr. Richard Pew
Bolt Beranek & Newman
50 Moulton St.
Cambridge, MA. 02138

Dr. Fred Muckler
NPRDC
San Diego, CA. 92152

Dr. Jesse Orlansky
IDA
400 Army-Navy Drive
Arlington, VA. 22209

Dr. Warren Teichner
Dept. of Psychology
New Mexico State University
Las Cruces, NM.88003

Dr. Alfred F. Smode
NTEC (OOT)
Director, TEAG
Orlando, FL.32813

Dr. Craig Fields
DARPA
1400 Wilson Blvd.
Arlington, VA.22209

Dr. Marvin Denikoff
Office of Naval Research
800 N. Quincy St.
Arlington, VA. 22217

Col. Henry L. Taylor
ODDR&E
Washington, DC 20301

Cdr. Charles Teisen
Human Factors Engineering
NADC
Warminster, PA. 18974

Dr. Robert G. Smith, Jr.
CNO (OP-9874)
Pentagon 5D772
Washington, DC 20330

Dr. John A. Modrick
Human Factors Branch
Honeywell, Inc.
2600 Ridgeway Parkway
Minneapolis, MN.55413

Mr. Jim Baker
ARI
1300 Wilson Blvd.
Arlington, VA. 22209

Dr. Earl A. Alluisi
Dept. of Psychology
Old Dominion University
Norfolk, VA. 23508

Dr. Robert H. Ahlers
NTEC (N-215)
Orlando, FL. 32813

Dr. Tom Nyman
ODDR&E/OAD (E&PS)
Washington, DC 20301

Dr. Joe Zeidner
ARI
1300 Wilson Blvd.
Arlington, VA. 22209

Dr. Edgar M. Johnson
ARI
1300 Wilson Blvd.
Arlington, VA. 22209

HQ AFSC/SGB (Lt. Col. Trumbo)
Andrews AFB, MD. 20334

Dr. Roger L. Grossel
OASD (I&L)
Washington, DC 20330

Maj. Gen. Jasper A. Welch
ACOS/Studies & Analysis
HQ USAF
Washington, DC 20330

Lt. Gen. W. L. Creech
ESD/CC
Hanscom AFB, MA. 01730

Dr. James G. Miller
President
University of Louisville
Louisville, KY.

Dr. S. Zacks
Chairman, Dept. of Math & Statistics
Case Western Reserve University
Cleveland, OH.

Dr. Shanti S. Gupta
Dept. of Statistics
Purdue University
Lafayette, IN.

Dr. James R. Thompson
Dept. of Mathematical Sciences
Rice University
Houston, TX.

Maj. James A. Velezis
HQ DA (DAMA-CSS)
Washington, DC 20301

Dr. Robert S. Andrews
ARI
1300 Wilson Blvd.
Arlington, VA. 22209

Dr. Clyde Coombs
Dept. of Psychology
University of Michigan
Ann Arbor, MI.

Dr. Robert M. Thrall
Dept. of Mathematical Sciences
Rice University
Houston, TX.

Dr. Seth Bonder
Vector Research Inc.
Ann Arbor, MI.

Dr. G. Bracken
IDA
400 Army-Navy Drive
Arlington, VA. 22209

Lt. Col. R. A. Geesey
HQ USAF/SAGR
Washington, DC 20301

Dr. Robert Kalaba
Dept. of Economics & Biomedical Engr.
University of Southern California
Los Angeles, CA.

Col. R. H. Stuart
HQ USAF/SAGR
Washington, DC 20301

Dr. Thomas Hallman
Dept. of Mathematics
University of Georgia
Athens, GA.

Dr. H. O. Hartley
Institute of Statistics
Texas A&M University
College Station, TX.

Dr. Thomas Sasty
Warton School
University of Pennsylvania
Philadelphia, PA.

Dr. Robert L. Lavner
Mathematics Division
US Army Research Office
Durham, NC.

Dr. Julien M. Christensen
Dept. of Industrial Engr. & Ops. Res.
Wayne State University
Detroit, MI. 48202