Ames Grant
IN-53-CR
211802
p 300

THE OHIO STATE UNIVERSITY

The Organization of Perception And Action in Complex Control Skills

Richard A. Miller
Department of Industrial and Systems Engineering

Richard J. Jagacinski
Department of Psychology

National Aeronautics and Space Administration
Ames Research Center
Moffet Field, California 94035

Grant No. NAG 2-195
Final Report

May 1989

THE OHIO STATE UNIVERSITY

The Organization of Perception And Action in Complex Control Skills

Richard A. Miller
Department of Industrial and Systems Engineering

Richard J. Jagacinski
Department of Psychology

National Aeronautics and Space Administration
Ames Research Center
Moffet Field, California 94035

Grant No. NAG 2-195
Final Report
RF Project 763264/714826

May 1989

ACKNOWLEDGEMENTS

This document constitutes the final report for Grant NAG 2-195 and it is basically the dissertation of Alex Kirlek. The work also had the benefit of the involvement of several additional dedicated graduate students including Lisa Messing, Lynn Lytton, and Brian Plamondon. Their individual work has been reported in previously submitted documents. All involved with the project gratefully acknowledge the encouragement and support of E. James Hartzell, project monitor.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Problem Statement

This work is an attempt to describe the perceptual, cognitive, and action processes that account for highly skilled human performance in complex task environments. In order to study such performance in a controlled setting, a laboratory task was constructed and three experiments were performed using human subjects. Then, a general framework was developed for describing the organization of perceptual, cognitive, and action processes. This framework is intended to apply not only to the laboratory task but to a hopefully much wider range of task environments.

The laboratory task was a simulation that required human crews to perform manual and supervisory control of a fleet of vehicles. The term "supervisory control" stems from the analogy between a supervisor of a subordinate staff in an organization of people and the human controller of a modern, computer-based, semi-automatic control system (Sheridan, 1987). In such systems, the human operator is typically described as performing many knowledge intensive functions such as prediction, planning, and decision-making.

At a minimum, the operator is responsible for directing the activity of the automated systems, monitoring the status of the controlled process, and intervening in the case of unexpected environmental disturbances or equipment malfunctions. Examples of such systems include semi-automated manufacturing systems, air traffic control systems, and large electricity generation plants.

While the laboratory task used here was quite complex, especially when contrasted with most other experimental tasks, it was of course far simpler than the example control systems mentioned above. One major dissimilarity between the simulated task and actual supervisory control tasks was that the former did not require crews to detect, diagnose, and repair failed system components. As a result, the human cognitive processes involved in dealing with novel or rare events were not studied in this research. On the other hand, the laboratory task did preserve many other properties of actual complex control tasks. Perhaps most important were those requiring crews to decompose the overall task goal into a set of subgoals, to constrain and coordinate the many degrees of freedom available for system control, to rapidly process a large set of graphically displayed information, and to cope with simultaneous task demands.

What made the study of human behavior in the laboratory task particularly challenging from a modeling perspective was the breadth of issues that had to be considered to provide a comprehensive description of the psychological processes underlying skilled performance. Much psychological research has relied upon decompositions of phenomena into discrete categories, such as perception, cognition, and motor control. There is little doubt that this has been a productive research strategy, at least from a methodological perspective. But from a theoretical perspective, the possibility exists that decomposing psychological processing into these discrete problem areas may be counterproductive. As Pylyshyn (1982) has noted,

> [T]his is not to imply that we cannot study problems in these areas, but rather to suggest that categorizing them in such terms may not reveal (indeed, may cover up) the way in which they individually and collectively contribute to producing intelligent behavior. (p. 70)

As Pylyshyn's comment suggests, it sometimes appears as if too little attention is paid to the task of theorizing how the collection of these distinct psychological modules are interfaced and coordinated to yield goal directed behavior in complex, dynamic environments.

In this research, though, the issues of system-level organization and coordination among these modules occupied center stage. Two factors prompted this research focus. The first concerned the need to allow for the theoretical possibility that skilled behavior in the laboratory task was due as much to efficiencies arising from the organization of psychological processes as it was due to the presence of individually efficient processes. In other words, it could be possible that much of what accounts for high level skill might involve the way in which the information processing demands of a complex task are effectively decomposed into separate tasks for the perceptual, cognitive, and motor systems. Given that each of these systems might have its own "preferred" mode of information processing, it would appear at least possible that skill development might involve obtaining an efficient division of labor among these systems so that each system performs operations suitable to its preferred processing style. To at least allow for this possibility, it was imperative that issues of organization and coordination among these systems be given primary consideration.

The second reason that organizational issues were emphasized in this research concerns the methodological difficulties associated with the unobservability of psychological processes. Independent of the validity of the skill development hypothesis stated above, it is still mandatory that distinct functional roles for the perceptual, cognitive, and motor systems be identified for the purpose of psychological modeling. This problem is particularly acute with respect to cognitive processes, because neither the input nor output interfaces are anchored in observable phenomena. While the inputs to the perceptual processes and the outputs of the motor processes are typically available for empirical measurement, neither interface of the cognitive system is amenable to such study. No matter how tasks are constructed in an attempt to observe the behavior of cognitive processes, we find we are once again only directly observing the behavior of motor processes, whether they be arm movements or the production of speech. Similar problems are associated with trying to clearly identify the inputs to cognitive processes when only the inputs to perceptual processes can be observed.

These comments are not intended to be construed as a neo-behavioristic argument against the utility of postulating cognitive processes. Rather, all that is suggested is that it may be dangerous to make naive assumptions about the nature of the information that is contended to serve as the inputs and outputs of cognitive processes. The danger is due of course to the fact that if the assumptions concerning the input and output languages are false, there is no possible way for the characterization of cognitive processing to be correct. In addition, representational assumptions concerning cognitive processes also constrain theories of perception and action, since models of these processes then become obligated to communicate with cognitive processes in a certain format.

Perhaps, then, the greatest potential contribution of this research is the development of an overall framework for describing how perceptual, cognitive, and action systems might be organized and integrated to produce highly skilled performance in complex task environments. Unfortunately, the global nature of this research is also likely to be its greatest potential source of criticism. This point is well taken, since very few questions concerning the independent and detailed activity of perceptual, cognitive, or action mechanisms have been addressed with much depth. On the other hand, the work does address a set of novel issues that would appear to surface only when this discrete set of psychological mechanisms is viewed as an integrated whole operating within the context of complex human action. The modeling framework developed in this work is concerned primarily with these system level issues concerning the global organization of psychological processes.

Methodological Issues

In this work, the modeling framework played the role of a psychological theory for which a generative process model was developed and applied to describing the psychological processes of human crews performing the laboratory task. A generative model is one which uses similar information (inputs) and controls (outputs) to those used by human subjects to perform a task. A generative model typically produces a stream of behavior which can be compared to human behavior to assess the validity of the model. For example, a regression model that only indicates the statistical relations between task parameters and summary human performance measures would not be generative. A process model is one that is contended to describe, at some level, the actual time course of human information processing that occurs during task performance.

Process models, like the one constructed in this work, typically postulate the existence of directly unobservable psychological mechanisms and processes. Within cognitive psychology, this approach to describing behavior has its origins at least as early as the work of Miller, Galanter, and Pribram (1960). With their TOTE unit mechanism for describing plans, they were among the first researchers to break with behavioristic tradition and employ theoretical constructs that mediated between stimulus and response. With the adoption of the "information processing" metaphor (Neisser, 1967), cognitive psychologists began to postulate theoretical constructs of increasing complexity to account for behavior that did not appear to be easily described within the language of behaviorism.

Although these information processing models appeared to offer the theoretical resources to describe a wider range of behavior than could the behavioristic models, some have faulted this type of theorizing for its vagueness. For example, Neisser (1976) criticized the proliferation of ill-defined "boxes" and "arrows" within many of these models. He pointed out that models constructed by simply decomposing an information processing task into a set of presumably simpler processors (boxes) connected by paths (arrows) of information flow left many critical questions unanswered. Little consideration was given to the mechanisms by which the processors operated or to specifying the exact nature of the information that was assumed to flow between them.

An explicit focus on <u>mechanism</u>, on the other hand, has been one of the major concerns of researchers operating within the "symbol manipulation" metaphor (Newell and Simon, 1972). While Newell and Simon were far from being the only researchers concerned with being explicit about mechanisms, they were among the first to make this concern an active constraint on their theories and models of cognition. By their adherence to the constraint that psychological models should be implemented as computer simulations to ensure completeness and explicitness, these researchers were highly influential in lessening the scientific acceptability of vaguely stated models of the type criticized by Neisser.

Newell and Simon's research has had a large impact on psychological methodology in other ways as well. Newell (1973) has made the observation that the purely "bottom-up" or data-driven approach in psychology has, as of yet, mainly produced a discrete set of laboratory curiosities of questionable significance. He then forcefully argues that a "top-down", or theory-driven approach offers the best hope for advancing the field. Although actual scientific progress would seem to rely upon a rich interplay between these two approaches, Newell appears to have been influential in at least tilting the balance toward a top-down approach to cognitive psychology research. Of course, it is still too early to tell with certainty whether or not this strategy will be scientifically rewarding. On the other hand, a top-down strategy appears to be almost a methodological necessity for structuring the approach to investigating rich behavioral situations such as the one studied in this research.

This is not to say, of course, that by using a top-down methodology in this research the scientific obligation to empirical adequacy is ignored. Rather, one of the primary reasons for using generative models is to provide a method for indirectly measuring the properties of the theorized psychological mechanisms. Such measurements are performed by using the model to assess the degree to which hypotheses about the mechanisms' properties are in agreement with observable behavior. Specifically, the behavior of the model parameterized in accordance with a set of hypotheses is compared to human behavior. If the parameterized model produces behavior in agreement with human behavior, the sufficiency of the set of psychological hypotheses is established.

The sufficiency demonstration in this work was performed by implementing the model as a computer program that was capable of performing the task at the performance level of the most highly skilled human crews. By a process of parameter manipulation, the model was then used to attempt to mimic the behavior of three human crews. A set of performance measures were constructed to provide detailed comparisons between model and crew performance. The model was able to generate behavior in approximate agreement with each of the three crews. In addition, the model constructs provided for a reasonably concise description of inter-crew differences.

Given that the computer program was able to achieve approximate behavioral validity, an important question becomes determining what psychological significance should be attached to this sufficiency demonstration. One aspect of this question concerns the breadth of the possible psychological implications, and its answer lies in properly identifying those aspects of the modeling framework that are task-independent. For this reason, the framework is presented in a task-independent fashion before it is applied to describing performance in the laboratory task.

A second component of the question of significance concerns the appropriate psychological interpretation of the behavior of the computer program. To answer this question, it is critical that the distinction between the psychological theory and its model as a computer program be kept clear. This distinction must be maintained so that confusions will not occur regarding the intended relationship between the operation of the computer program and the operation of psychological processes.

The psychological theory which is developed in this work is a set of statements or assumptions about the nature of possible psychological processes operative during human performance in the laboratory task. Under the standard model theoretic definition, a model of a theory is a structure in which the statements of the theory can be interpreted as true. The computer program is one such model, for it has been designed to operate in agreement with the collection of statements comprising the theory. Computer programs expressed in different languages and other adequately chosen structures may also serve as models of the same theory. The reason it is important to keep the distinction between the psychological theory and the computer program clear is that there exist an additional collection of statements that are true of the program but are <u>independent</u> of the theory. These statements concern implementation details which are specific to the choice of the model for the theory, but are not a part of the theory itself. It is crucial that these implementation details should not be erroneously construed as making claims about psychological processes.

While the distinction between the substantive claims of the theory and the details of computer implementation is clear enough in principle, it can be quite difficult to make this distinction in practice. The design of the computer model is a product of both the top-down constraints imposed by the psychological theory and the bottom-up constraints imposed by the structure of the computer hardware and software. It can be quite difficult to determine whether a particular section of computer code should be interpreted literally as simulating the operation of a hypothesized psychological mechanism, or on the other hand, whether that code should be construed as a somewhat arbitrary implementation that only preserves the gross functional (input-output) characteristics of a psychological mechanism.

In order to reduce the possibility of such confusions, the modeling components of this research are presented in an incremental fashion. The presentation begins with a discussion of the general modeling framework, followed by an application of this framework to modeling performance in the laboratory task, and ending with an explicit description of the operation of the task-specific mechanisms that is closely tied to the computer implementation of the model. This decomposition of the work and other issues related to presentation are described in the last section in this chapter.

Overview of Results

As mentioned above, the computer model constructed in this research was able to perform the laboratory task at the level of the most highly skilled human crews. In addition, by manipulating model parameters to mimic the behavior of three different crews, the model could be used to provide a reasonably concise description of inter-crew differences. The computer model was an implementation of a psychological framework developed to describe how a complex information processing task might be decomposed into separate tasks for the perceptual, cognitive, and action systems. Although the modeling exercise will be discussed in detail in later chapters, some of the more interesting results of this work will be discussed here.

One class of results concerns the psychological assumptions that were made concerning the decomposition of the information processing task into functions for the perceptual, cognitive, and action systems. Since these assumptions were shown to be sufficient for the generation of skilled behavior, they should at this point be considered candidate descriptions of the mechanisms underlying complex control skills. Due to the problems of unobservability discussed above, the assumptions were generated in a data-driven fashion where possible. Since at least the displayed information and observable control activity could be identified, the approach that was used to identify distinct functional roles for the perceptual, cognitive, and action systems was to work from the "outside-in".

Working in from the action side, the issues that were addressed concerned hypothesizing how the development of highly structured action mechanisms could possibly simplify the functional role of the cognitive system by constraining and defining its effective outputs. Working in from perception, consideration was given to the problem of understanding how the development of highly structured perceptual mechanisms might simplify the role of the cognitive system by constraining and defining its effective inputs. Only after these task-specific perceptual and action mechanisms were defined could questions concerning cognitive system be addressed.

By approaching the problem in this fashion, the view of psychological processes in skilled performance that emerged in this research was one of heavy reliance upon highly structured, task-specific perceptual and action systems. The high degree of complexity present in these peripheral systems was described as allowing for much of the burden for the production of complex behavior to be removed from the central, or cognitive system. The role played by the perceptual systems in this regard was suggested to derive from their ability to abstract information from the displayed environment that was highly relevant to the task of action selection.

Specifically, the perceptual systems were described within the framework of ecological perception as mechanisms devoted to the detection of action-oriented information, or "affordances" (Gibson, 1966). Assuming a sensitivity to affordances allowed for assumptions to be made concerning how perceptual systems might categorize and differentiate the displayed environment along dimensions that were efficient for the task of action selection. With affordance derived categorization, efficiencies resulted from the fact that many apparently distinct environmental situations could be treated in an identical way by the cognitive system. With affordance derived differentiation, efficiencies resulted from the fact that many apparently similar environmental situations could be beneficially treated as distinct by the cognitive system.

Another type of efficiency that was found to derive from action-oriented perception was that the distinction between supposedly distinct types of cognitive tasks could be collapsed. This allowed for the assumption of a single cognitive mechanism to perform a major component of all action selection tasks, independently of whether these tasks would traditionally be described as involving decision-making, coordination, or planning. All three of these types of tasks shared a single cognitive mechanism, the operation of which was quite simple. Namely, this mechanism was assumed to select the action with the highest affordance value, subject to a single constraint that resolved competition between proposed actions.

Whether behavior typically described as involving decision-making, coordination, or planning was generated by this mechanism was determined by the affordance information upon which the mechanism operated. Decision-making was produced by allowing the cognitive mechanism to operate upon the current action affordances for each of the independent craft under crew control. Coordination among the craft was produced by applying the same cognitive mechanism to a different set of affordances. These affordances were based on a perceptual sensitivity to higher-order relations among the multiple craft. These relations indicated opportunities for potential beneficial cooperation among the craft. In this way, the distinction between decision-making and coordination could be assumed to be primarily perceptual, rather than cognitive in nature.

For planning the future activities of the craft, once again the same cognitive mechanism was used. For planning, though, the cognitive mechanism was assumed to operate upon the affordances that were predicted to exist in the environment at a future time. Thus, the planning process was more complex than the decision-making and coordination processes since an ability to predict future environmental states had to be assumed. On the other hand, the planning process was still quite efficient since it used the same cognitive and perceptual mechanisms as were used for the decision-making and coordination tasks. Additional perceptual mechanisms were not required for planning because the original set of perceptual mechanisms were assumed to be able to operate upon the predicted, internally imaged environmental state.

Although this discussion has focused on the important role that was assumed to be played by perceptual systems, the modeling framework was also based on the assumption that highly structured action systems make a significant contribution to the production of skilled behavior. The action systems were assumed to operate as a hierarchically structured control system. This arrangement of action mechanisms was assumed to be responsible for the generation of routinized control activity. In such systems, the lower-level mechanisms possess a degree of control autonomy so that much of the burden for the production of complex behavior can be relieved from the higher-level mechanisms. In this way, the assumption of highly structured action systems helps to allow for the assumption of relatively simple cognitive processing mechanisms.

This brief overview of the modeling results has left out many of the interesting empirical findings concerning crew performance and inter-crew differences. It is hoped, though, that this discussion has provided enough of an introduction to the work to encourage the reader to explore both the theoretical and empirical issues in greater detail.

Document Organization

Following a description of the laboratory task and experiments in Chapter II, a performance characterization of each of the human crews are developed in Chapter III. Readers primarily interested in the empirical aspects of human performance in the laboratory task may then skip to the discussion of the conclusions of the modeling results in Chapter VI. This discussion provides an interpretation and explanation of the individual differences in human performance that were observed in the experiments.

Chapter IV discusses the general theoretical framework that was developed for describing the possible organization of perceptual, cognitive, and action mechanisms accounting for skilled performance. This framework is independent of the specifics of the laboratory task and is intended to apply to a wide range of task environments. Results from the literature are reviewed here to motivate the present approach. Readers, though, who are primarily interested in the design and operation of the task-specific model may choose to omit this chapter and move to Chapter V where the task-specific constructs are discussed. A familiarity with the issues raised in Chapter IV is not crucial to understanding the design of the task-specific model, although readers who find the model's design to be arbitrary or unconventional may find a discussion of these issues to be beneficial.

The most specific description of the model is given in Chapter VI. Here, the model parameters are identified and the way in which they were manipulated for describing individual human crews is discussed. This chapter describes the equations and logically based routines that were used to implement the model as a computer program. This level of description adds no meaningful psychological content to the modeling approach. Rather, these equations and routines are relevant only in that they allow for the computer-based implementation of the hypothesized psychological constructs. The psychological meaning of the equations and routines is restricted to what they inherit by virtue of their functional role as somewhat arbitrary implementations of these meaningful constructs. For this reason, this chapter may be bypassed by readers whose primary interest is in the novel psychological contributions of this work.

The results of the modeling effort are discussed in Chapter VII. In this chapter, the behavior of each of the three human crews is interpreted in terms of the model constructs. Finally, research conclusions are presented in Chapter VIII. These conclusions include observations concerning the psychological relevance of this work and concludes with a discussion of how this work might be extended to provide a more comprehensive framework for describing the operation of human perceptual, cognitive, and action mechanisms engaged in skilled performance.

CHAPTER II

TASK AND EXPERIMENTS

Introduction

In this chapter, the simulated supervisory control task and the three experiments that were performed using human subjects will be described. Perhaps a word of caution is necessary concerning use of the term "experiment" in this work. As will be clearly seen in the modeling components of this work, this research does not follow the traditional pattern of hypothesis formulation and subsequent testing via experiment that is commonplace in experimental psychology. Although a few independent variables were identified and systematically manipulated (crew size, e.g.), the experiments must be seen as both a source of hypotheses as well as providing a means for assessing their validity. What makes this a defensible methodology is the distance between the level of description at which the empirical data are available and the level at which the psychological theories, mechanisms, and hypothesis are described. The empirical data exist in terms button pushes, key presses, and joystick manipulations. The psychological theories and models used to describe this behavior include constructs such as planning horizon, environmental affordance, and other directly unobservable entities. The distance between these two levels of description makes the task of inventing suitable high-level constructs, mechanisms, and methods for measurement of their properties a necessary prerequisite to formulating and testing hypotheses in terms of these constructs.

There is still some danger when using such a strategy of coming up with psychological mechanisms that are profoundly ad hoc, idiosyncratic, and at best, good for understanding human behavior only in the task from which their existence was induced. Two lines of defense are offered against this potential criticism. First, the psychological theorizing was strongly constrained by an array of other empirical data concerning human cognition. This fact makes it more likely that the constructs used here may have a wider domain of applicability than for modeling this task alone. A second defense against this criticism is to accept the task-specific nature of the model, but deny that this is really such a bad result. Perhaps a detailed account of human behavior in even this one complex task can provide insights that could not be gained from accounts with much greater range but possibly less detail. It is also the case that a clear and empirically grounded distinction between task-dependent and task-independent cognitive mechanisms and processes does not yet exist. Therefore, an a priori criticism of strongly task-dependent modeling may be premature and may be subsequently found to be the somewhat harmful result of a naive search for simple, general purpose mechanisms, where highly task-dependent and idiosyncratic cognitive processes may be the rule (Claxton, 1988; Allport, 1980).

Task Description

The empirical portion of this research is based on a laboratory simulation of a supervisory control task. The interface could be configured for either one- or two-person crew experiments. The work station for the one-person crew configuration is shown in Figure 1. The two-person configuration is similar although a subject is seated in front of each of the two large displays, and two button panels are used. The task required subjects to earn points by guiding five "friendly" craft through a forested 100 square mile "world" in search of "enemy" craft to be destroyed and cargo which were to be loaded and taken to home base to be unloaded. Points were awarded for destroying enemy craft and unloading cargo at home base. Points were deducted for the destruction of friendly craft (due to attack from enemy craft or by collision) and by having friendly craft run out of fuel. Each session lasted thirty minutes although it could be prematurely terminated if the main friendly craft, called the scout, was destroyed by an enemy craft or by

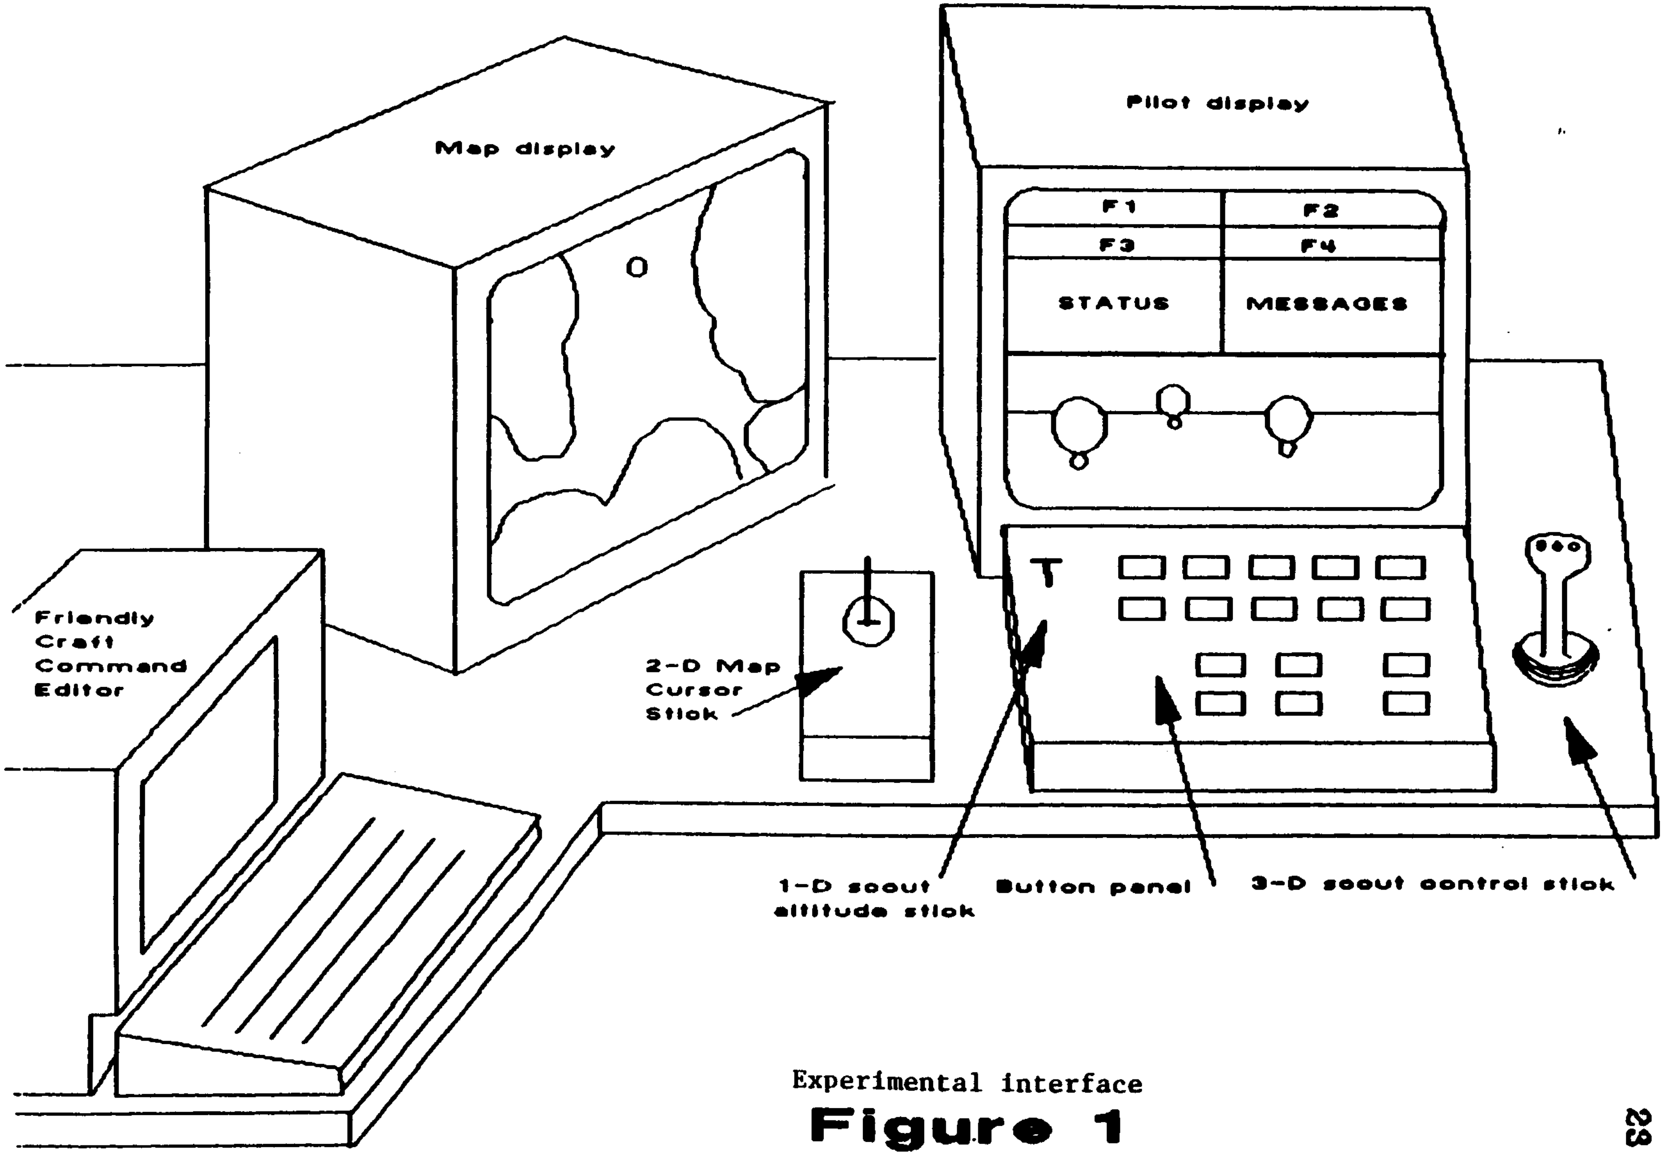

Experimental interface

Figure 1

colliding with another friendly craft or with a tree.

The interface was intended to represent the cockpit of the scout craft occupied by the subjects. Subjects had manual control over the scout via joysticks and push-buttons. A three-dimensional rate control joystick allowed for lateral, longitudinal, and yaw control of the scout. The manual control task was simplified by excluding pitch and roll from the simulated helicopter dynamics. In addition, the dynamics that were used resulted in inherent stability for the scout. Therefore, when subjects diverted attention from the manual control task the scout coasted to a rest at fixed altitude. A one-dimensional rate control joystick was used for altitude control of the scout. Subjects could also choose to use automatic controls for these purposes. Two buttons could be used to automatically transition the scout to a fixed high or low altitude (above or below tree level) in lieu of the manual altitude control. An autopilot with a tree avoidance mechanism could be used instead of the manual control joystick, although autopilot control only operated at 75% maximum ground speed with double normal fuel consumption.

In addition to the scout craft in which the crew resides, subjects also had control of four additional semi-automated friendly craft. These craft were functionally similar to the scout craft although they could not be manually controlled. Instead, subjects were required to construct lists of action commands for the friendly craft on a text editor specifically designed for the simulation. A cursor on the map display (a top-down display of the entire 100 square mile world) could be used to send the friendly craft to specified locations in the world and to enter scout autopilot waypoints. The friendly craft actions which were used to construct action strings include: "G" - goto waypoint specified by the map cursor location; "A" - attack an enemy craft with a missile; "L" - load cargo; "U"- unload cargo and replenish fuel and missiles at home base; "P" - patrol in a circular pattern; "^" - rise above tree level; "v" - descend below tree level. Subjects could construct strings of these actions prior to the time at which their execution was desired. This feature was included in an attempt to make some components of planning behavior directly observable. These command strings could be modified in real time by a series of push-buttons that allowed subjects to abort current strings, enable new strings, interrupt active strings, and skip single actions within active strings.

Twelve cargo and eighteen enemy craft were distributed throughout the world and were not shown on the map display until they were sighted by sensors centered at the scout and each of the four friendly craft. The density of cargo in the lightly forested regions was approximately twice the cargo density in the open and heavily forested regions.

Actual cargo locations were determined by generating locations with a pseudo-random number generator according to the forest density constraints.

The sensors used by the five craft to sight cargo and enemy craft will be informally referred to as "radar", although explicit radar properties were not simulated. The radar had a range of 1.5 miles while the scout was above tree level, and a range of 0.4 miles for the scout below tree level and for the friendly craft at all altitudes. The eighteen enemy craft consisted of six enemy helicopters that moved at the same speed as the friendly craft, six tanks that moved slightly slower than the friendly craft, and six stationary emplacements. The locations of the stationary emplacements were generated with the same method used to place the twelve cargo. The initial locations of the mobile enemy craft were random, and the enemy craft motions were determined with a semi-random walk algorithm. The motions were only semi-random in that each one of the six tanks and each of the six helicopters were constrained to occupy one of six rectangular sectors into which the world was divided. This feature kept the mobile enemy craft dispersed within the world. An enemy craft could only exit its designated sector if it was attacking a friendly craft.

The primary task involved searching the world with radar to discover and process the twelve cargo and eighteen enemy craft. A major component of the task involved management of the scout and friendly craft resources. All five craft began the mission with a full tank of fuel and five missiles that could be used to attack enemy craft. Fuel consumption rates were designed so that each craft was required to return to home base to refuel at least once at some point in the middle third of the mission. Each enemy attack episode typically used between one and three missiles. An unarmed friendly craft that engaged an enemy craft from which it could not escape (an enemy helicopter) was destroyed. The scout helicopter and friendly helicopters could not escape from enemy helicopters because all the helicopters traveled at the same maximum speed. The scout and friendly craft could, though, escape from enemy tanks as the tanks traveled at 75% of the maximum speed of the helicopters.

Friendly craft could be resupplied with missiles at any time by returning to home base. The number of cargo that could be carried by a friendly craft at any one time was constrained and thus represented another complication for resource management. Although this number depended on the (random) weight of the loaded cargo, the maximum number of cargo that could typically be carried was three. The problem of achieving an efficient utilization of craft resources required behavior that would typically be described as requiring extensive planning and anticipation.

Subjects used two dynamic, color, computer-generated information displays to perform the task. The map display, the center display in Figure 1, showed irregularly shaped regions of open ground, light forest, and heavy forest in different colors (brown, light green, and dark green, respectively). Home base was displayed as a white circle, and the positions of the scout and four friendly craft were displayed as small, numbered, blue circles. Any enemy craft and cargo that had been discovered were also displayed as small color coded circles. The positions of moving objects were updated once per second.

A second display, the rightmost display in Figure 1, was divided into three distinct areas. The bottom portion of this display showed a dynamic inside-out view of the terrain through which the scout was traveling. When the scout was in an open region, zero to three trees were typically within the 2000-foot scout viewing range and appeared on the inside-out display. When the scout was in heavy forest, eight to thirty trees appeared, and the resulting tree avoidance task was considerably more difficult. The middle portion of this display indicated resource information, warning messages (for fuel, lock-on by enemy radar, e.g.), mission time remaining, the state of scout control modes (manual vs. automatic), and points earned. Resource information for the scout (fuel, altitude, number of missiles remaining, weight capacity remaining) was the default information shown in the resource area, although corresponding information for each of the four friendly craft was available on a call-up basis. The top portion of this display indicated the lists of commands that had been assigned to each of the friendly craft.

In the two-person crew condition, one subject, called the navigator, was primarily responsible for the control of the four friendly craft via the text editor and push-buttons. The other subject, called the pilot, was primarily responsible for the control of the scout craft via manual control joysticks. The pilot was provided with a duplicate set of push-buttons to aid the navigator in the control of the friendly craft although such activity was rarely observed in the experiments described below.

In the one-person condition, the subject was required to time-share between the task demands of friendly craft control via the text editor and control of the scout. Autopilot control of the scout was available to relieve the subject from the demands of manual scout control but at the cost of the significant resource penalties mentioned above.

Experiments

Three experiments were performed by three different experimenters. The first two experiments (Plamondon, 1985; Lytton, 1987) were identical except for the manipulation of crew size. Each of these experiments used five subjects (or two-person teams) for twenty sessions. Four different forest configurations and four different placements of the cargo and enemy craft were used so that subjects would not be able to remember the locations of these objects from session to session. Two different point payoff structures were used to vary whether cargo loading and unloading or enemy craft destruction was emphasized each mission. The maximum points attainable for each mission, though, was held constant across payoff structure. Details of these experiments and a descriptive characterization of performance differences due to crew size can be found in Lytton (1987) and Miller, Jagacinski, Plamondon, Lytton, and Kirlik (1987). The task instructions given to subjects are provided in Appendix A.

One major result of the two initial experiments was that very large performance differences were observed between crews within the same crew size groups. Although the mean performance of the five one-person crews was about one third the mean two-person crew performance (758 vs. 2514 points), one of the one-person crews was able to score higher than the mean two person score. Due to these strong individual differences, the modeling approach used here attempts to describe behavior at the level of the individual subject, since aggregating data over multiple subjects could lead to misleading or meaningless information. Specifically, the data from these experiments that will be modeled are the highest scoring two-person crew and the highest scoring one-person crew. The highest scoring one-person crew will be referred to as Crew 1. The highest scoring two-person crew will be referred to as Crew 2.

Another implication of these strong individual differences for modeling skilled performance is that twenty sessions may not have been enough experience to allow all subjects to reach an asymptotic performance level. In order to better understand the nature of skilled one-person crew performance, an additional experiment was performed by the present author using three subjects for a larger number of sessions. One of these subjects (the present author) was highly skilled at the task and served as an "expert" subject and as a trainer for the other two subjects. This subject will be referred to as Crew E.

The only data from this last experiment considered here is the performance of the expert subject. As one goal of this experiment was to develop and define a consensual expert strategy for performing the task, the expert subject played the role of an on-line advisor to the two trainees for most of the duration of the experiment. At the end of the experiment, eight one-person crew sessions were performed by each of the trainees with limited interaction with the expert advisor. Although no explicit advice was given by the expert in these last eight sessions, verbal transcriptions indicated some interaction between the trainees and the expert so this data vas excluded from analysis.

To summarize, empirical data for three crevs vas chosen for use in the modeling exercise discussed belov: the best one-person and best tvo-person crevs from the first tvo experiments, and the expert subject from the last experiment. The final eight sessions for each crev vere used. These eight sessions comprise tvo sessions each on four different vorld configurations (forest location and object placement). One of the tvo sessions vith each vorld configuration vas performed vith the point structure payoff favoring loading and unloading cargo, vhile the other session vas performed vith the point payoff structure favoring attacking enemy craft. The entire experimental design is summarized in Table 1 on the following page.

TABLE 1

EXPERIMENTAL DESIGN

Design for Crews 1 and 2

Session	World Configuration	Payoff Emphasis
1	World 1	Enemy Craft
2	World 4	Cargo
3	World 1	Cargo
4	World 4	Enemy Craft
5	World 2	Enemy Craft
6	World 3	Cargo
7	World 2	Cargo
8	World 3	Enemy Craft

Design for Crew E

Session	World Configuration	Payoff Emphasis
1	World 2	Enemy Craft
2	World 3	Cargo
3	World 2	Cargo
4	World 3	Enemy Craft
5	World 1	Enemy Craft
6	World 4	Cargo
7	World 1	Cargo
8	World 4	Enemy Craft

Note: Sessions 1-8 are the final 8 sessions for all subjects

Payoff Structures

Event	Points Awarded with Emphasis on: Cargo	Enemy Craft
Cargo Unloaded at Home	400	100
Enemy Helicopter Destroyed	100	400
Enemy Tank Destroyed	60	240
Enemy Emplacement Destroyed	40	160
Friendly Craft Destroyed	-400	-400
Friendly out of Fuel	-100	-100
Maximum Possible Score	6000	6000

CHAPTER III

CREW PERFORMANCE PROFILES

Introduction

In this chapter, performance profiles of each of the three crews will be developed. The goal is to explore any performance differences between crews that should be captured by a generative descriptive model. First, a set of performance measures is described that were developed to aid in identifying strategic and competency related differences between the three crews. Next, the performance profiles of the three crews based on these measures are compared. Finally, inferences based on these profiles are made that suggest differences between crews in terms of the dimensions along which they have decomposed the task into subtasks, and the ways in which they have created and prioritized task subgoals. These inferences serve the role of hypotheses that will be subsequently evaluated in the modeling exercise which is discussed in the following chapters.

Performance Measures

Though it provides a general summary of proficiency, a scalar performance measure such as points scored is far too coarsely grained to provide meaningful diagnostic information in a task of the complexity studied in this research. Even in a relatively simple task such as single axis tracking, there are many distinct time histories the subject may produce and yet achieve the same scalar measure (e.g. rms error) of performance. This is especially true of sub-optimal performance, since optimal performance usually constrains the number of alternative strategies that can result in the theoretically maximum performance score. In the case of sub-optimal performance, the degrees of freedom available to the operator to control the system can usually be coordinated in numerous ways to achieve the same result.

This phenomenon may be summarized by noting that a given measure, in addition to providing a metric of performance, defines an equivalence class of behaviors that lead to each level of performance defined by that measure. Construed in this way, the task of developing measures that are suitably diagnostic is one and the same as determining those behavioral differences that are meaningful and those that are not, given the purposes of analysis. Meaningful behavioral differences for the purposes of generative modeling include those that are indicative of differences in the way crews have decomposed the task into subtasks and created and prioritized subgoals. Of particular interest is how the overall task demands were decomposed into separate demands for scout and friendly craft activity.

The performance measures have been designed to be sensitive to the alternative ways that crews can use the capabilities of the scout and friendly craft to search the world for cargo and enemy, process these objects, and manage craft resources. Measures include points scored, cargo and enemy craft sighted (friendly craft versus scout), cargo loaded (friendly craft versus scout), enemy craft destroyed (friendly craft versus scout), and amount of the world searched, also by friendly and scout. The time spent idle by the scout and friendly craft was also measured as an indicator of the inability of the operator to cope with the multiple task demands. The average time between loading a cargo and unloading it at home base was measured as a possible indicator of strategic differences between crews. The average number of cargo unloaded per trip to home base was measured as a possible indicator of whether crews were processing cargo serially or in parallel. (See Lytton, 1987; and Miller et. al., 1987 for an application of these measures to all five one and two-person crews.)

Figure 2 on the following page shows the points scored by each crew for each of the eight final sessions. As indicated in the diagram, each crew had one session with total points scored considerably lower than their average score over these eight sessions. For this reason, the lowest scoring session for each crew was discarded from the analysis. Comparisons were then made between each pair of

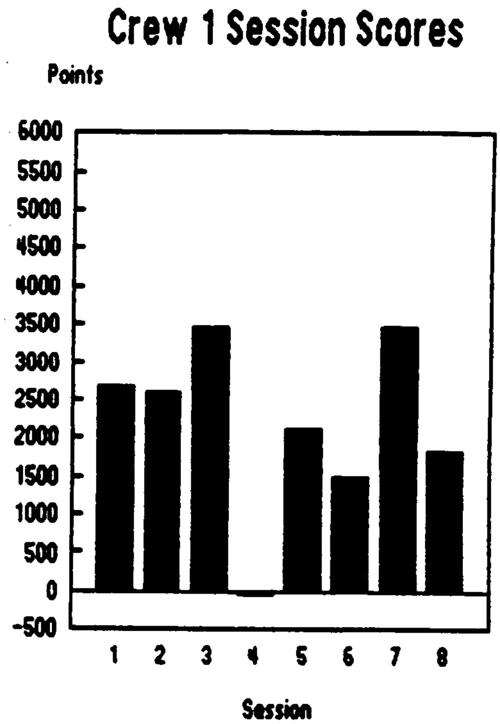

One Session Discarded Per Crew

Crew 1: Session 4

Crew 2: Session 2

Crew E: Session 6

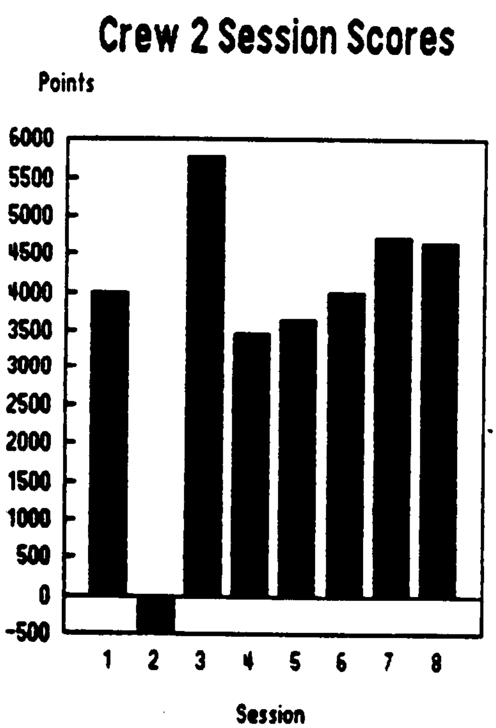

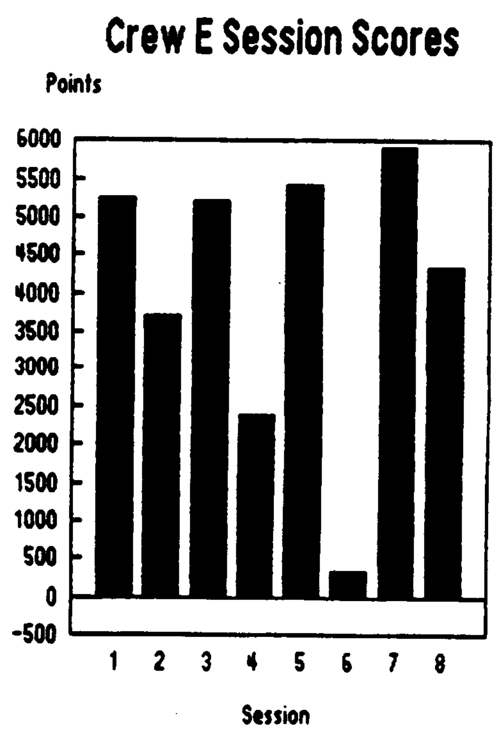

Crew performance showing outliers

Figure 2

crews for each performance measure over the seven remaining sessions using t-tests with alpha level = 0.05.

One potential criticism of this method of data analysis is that the probability of Type I error for the conjunction of tests is underestimated by probability of Type I error for the individual crew comparison tests. The method is defended on the grounds that both crew size (one versus two persons) and experience level (low versus high) should be individually treated at the 0.05 level in analogy with the way in which multiple experimental manipulations are treated in an analysis of variance (ANOVA). The present experimental design is essentially an incomplete 2 X 2 factorial design. It is incomplete since no observations are available in the high experience, 2-person crew condition. A standard ANOVA on this 2 X 2 factorial design would typically treat each of the two manipulations at the 0.05 alpha level, as is done here.

Comparison of Crew Performance

The average number of points scored per session appears in Figure 3 on the following page. The black bars in this and succeeding figures indicate values of the mean plus and minus one standard deviation. Crews 2 and E did not differ on this measure but both of these crews scored higher than Crew 1. Recall that Crew 1 indicates the highest scoring (excluding the expert subject) one-person crew, Crew 2

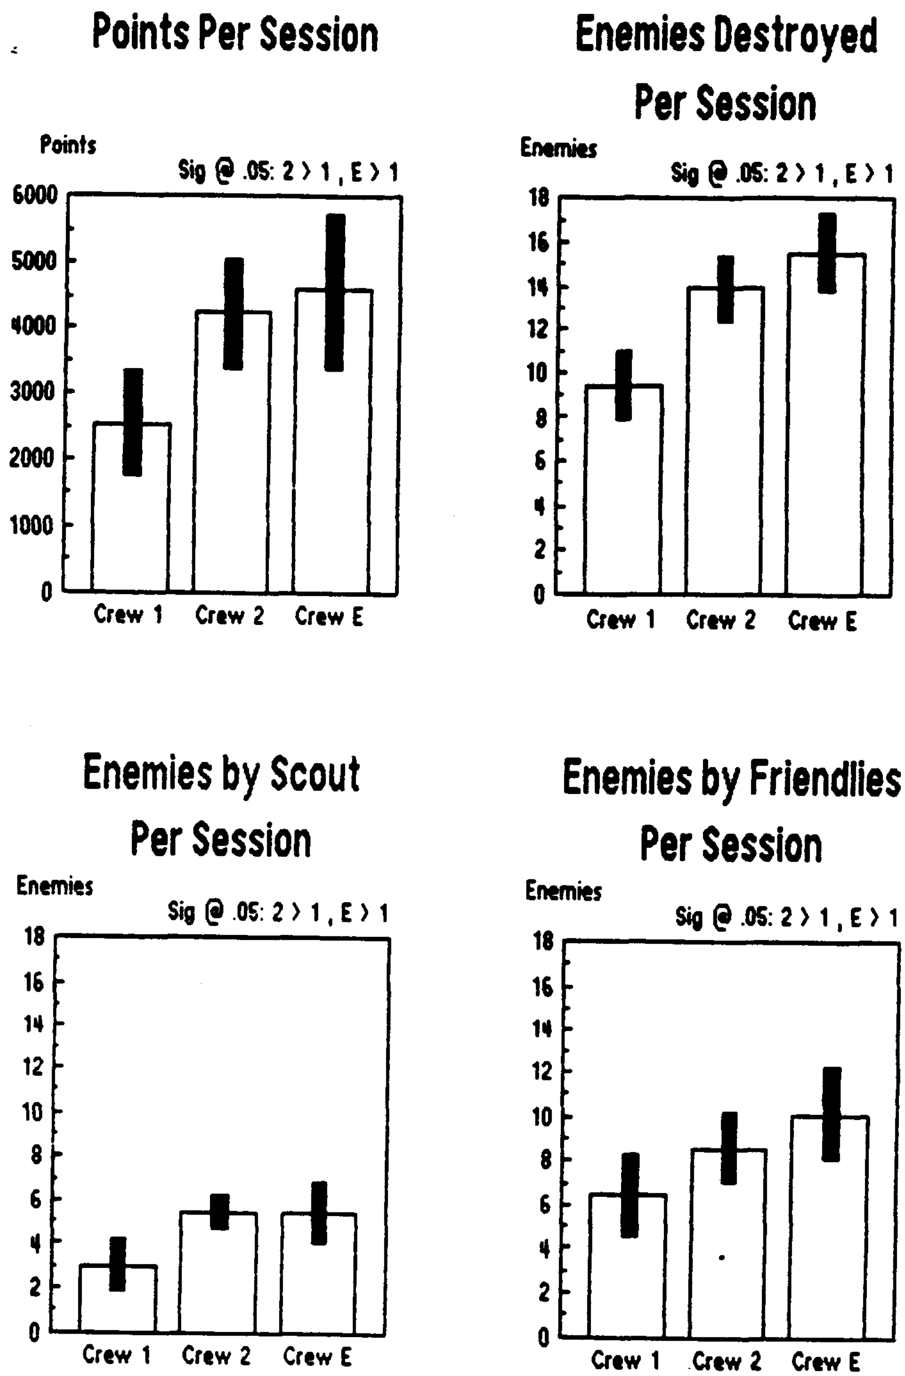

Crew comparisons of points and enemy craft destroyed

Figure 3

indicates the highest scoring two-person crew, and Crew E is the expert one-person crew. While it is not surprising that the two-person crew scored higher than the one-person crew, it is significant that the expert crew was able to score as well as the two-person crew.

It is clear that it must be explained how the expert crew, although time-sharing the demands of scout and friendly craft control, was able to score as well as the two-person crew who were able to decompose the task of scout and friendly craft control between two people. It should be noted that the time-sharing demands for Crew E were quite extensive, as this crew always used manual control of the scout. The other one-person crew, Crew 1, consistently used autopilot control of the scout, presumably as a way of coping with these time-sharing demands. Crew 2, the two-person crew with a member dedicated to scout control, consistently used scout manual control.

The next three graphs in Figure 3 indicate a similar performance trend for the three crews in terms of enemy craft that were destroyed. Whether measured in total or by scout and friendly craft independently, Crews E and 2 destroyed the same number of enemy craft, and both of these crews were superior to Crew 1. This measure, then, is no more diagnostic than points scored for explaining performance differences. Figure 4 on the following page suggests that the same trend holds for total cargo unloaded and cargo

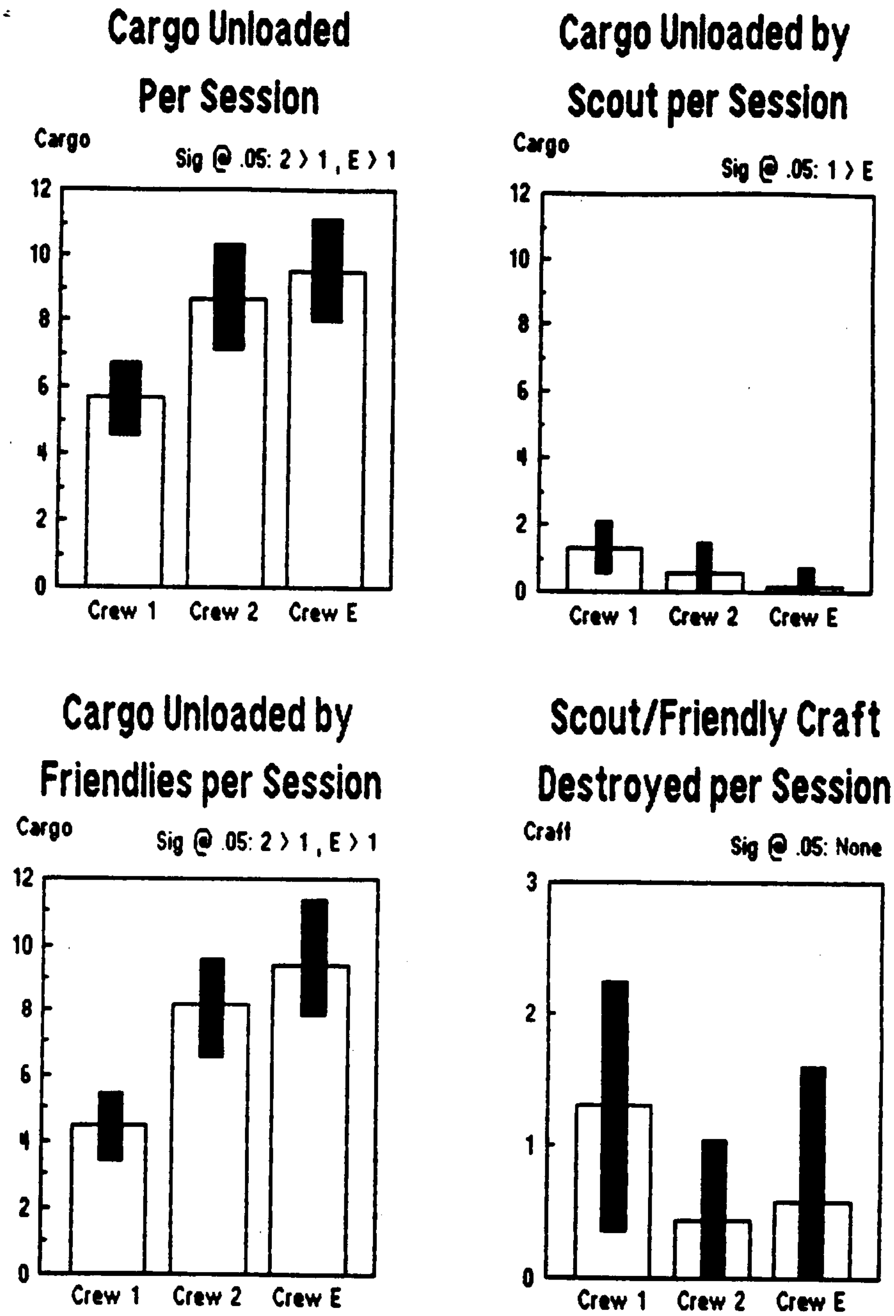

Crew comparisons of cargo processing and friendly craft destroyed

Figure 4

unloaded by friendly craft. For cargo unloaded by the scout, however, there was no difference between Crews 1 and 2. In addition, Crew 1 actually unloaded more cargo with the scout than did Crew E, although this difference is small in relation to Crew E's superiority to Crew 1 in terms of the number of cargo unloaded by the friendly craft. There was also no difference between any pair of crews in terms of number of friendly craft lost per session due to destruction by enemy craft.

The primary reason that Crew 1 unloaded fewer cargo than the other two subjects can be seen in the analysis of cargo discovery in Figure 5. Note that Crew 1 discovered less cargo with the scout than the other two crews, although he discovered the same number of cargo with the friendly craft as Crews 1 and E. All three crews discovered about 4.4 cargo per session with the friendly craft. Crew 1 unloaded an average of 4.3 cargo per session with the friendly craft, whereas the other two crews unloaded about twice this many with the friendly craft. Thus, Crew 1 appears to have been able to use the friendlies to unload all and only the cargo sighted by the friendlies, whereas the other crews were able to unload these cargo with the friendlies in addition to an equal number of cargo that were sighted by the scout. This fact suggests that Crew 1 rarely sighted cargo with the scout and then passed these cargo off to the friendlies to be loaded, whereas this was a common

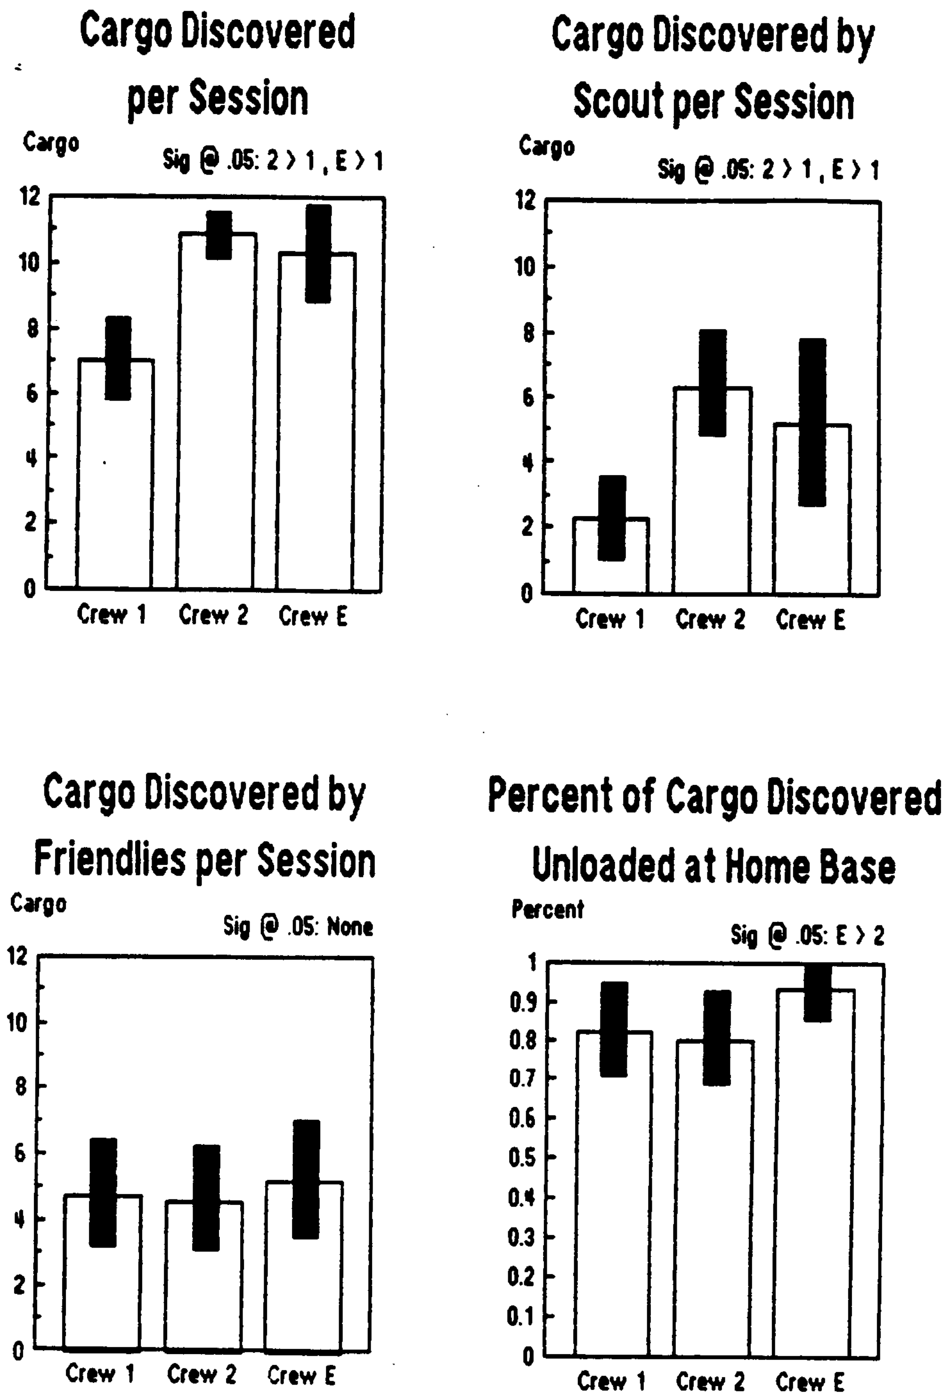

Crew comparisons of cargo discovery

Figure 5

occurrence for Crews 2 and E. This difference is most likely the result of the fact that, for Crew 1, the scout sighted so few cargo that the scout could be used to load all those that were sighted. For the other two crews, the scout was used much more effectively to sight cargo and thus cooperation with the friendly craft was required to load these cargo and carry them to home base.

Figure 5 also suggests the first difference between crews 2 and E, and this difference also appears to be due to cooperation between the friendly craft and the scout. Note that Crew E unloaded a greater proportion of the cargo that were discovered (within each mission) to home base than did Crew 2 (93% vs. 80%). One explanation of this result is that Crew E simply discovered the cargo earlier in the mission than Crew 2 (which would also result in a greater proportion unloaded at home), but an analysis of the cargo sighting times does not support this explanation. (No differences were found between any pair of crews on this measure.)

Figure 6 on the following page indicates that both Crews 2 and E unloaded the same number of cargo per trip home by the scout (0.4), and they unloaded the same number of cargo per trip home by each friendly craft (1.6). Thus, these two crews were similar in their differential usage of the scout and friendly craft. Both crews unloaded approximately four times as many cargo per trip home by a friendly

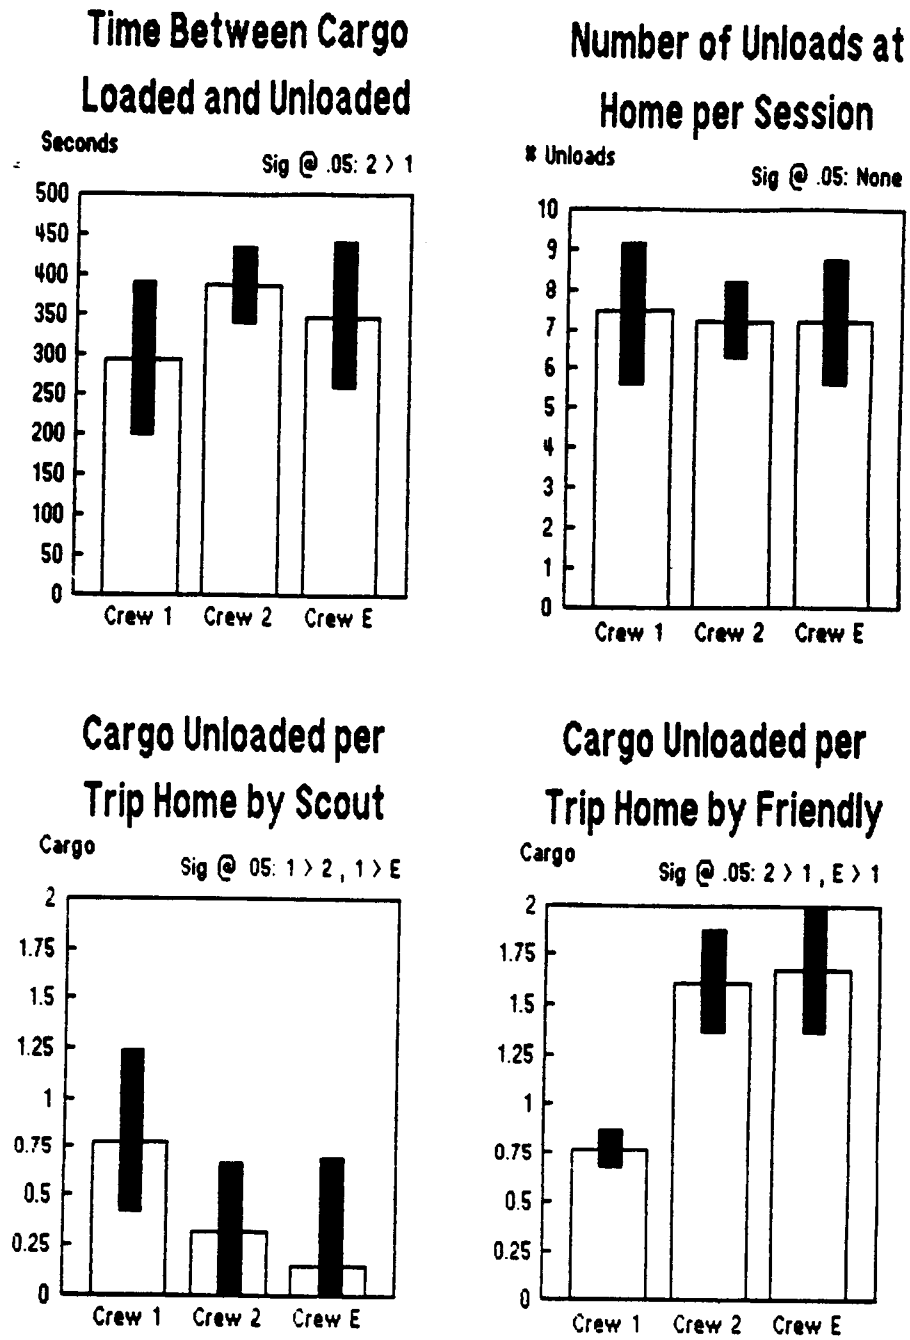

Crew comparisons of cargo processing

Figure 6

craft as were unloaded per trip home by the scout. On the other hand, Crew 1 showed no such differential usage of the scout and friendly craft. Both the scout and friendly craft unloaded approximately the same number of cargo per trip to home base for this crew. These facts suggest that one factor that may have contributed to Crew 1's poor performance was that he was insensitive to the functional differences between the scout and friendly craft. By essentially treating the scout as another friendly craft with regard to cargo processing, Crew 1 did not appear to make best use of the superior ability of the scout to discover cargo and enemy craft. Crews 2 and E, on the other hand, appeared to have recognized that the scout was best used to search the world due to its larger radar radius. These crews used the friendly craft to process the cargo that were discovered, instead of diverting the scout from its primary task of searching the world.

The next diagrams presented in Figure 7 suggest some additional differences between the three crews in terms of the utilization of the scout and friendly craft. A measurement of craft idleness was performed by calculating the percentage of session time that each craft was stationary. While this measure slightly overestimates true idleness due to stationary craft activity (cargo loading, e.g.), these stationary activities comprised a very small fraction of total mission time. While it is not surprising that Crew 2

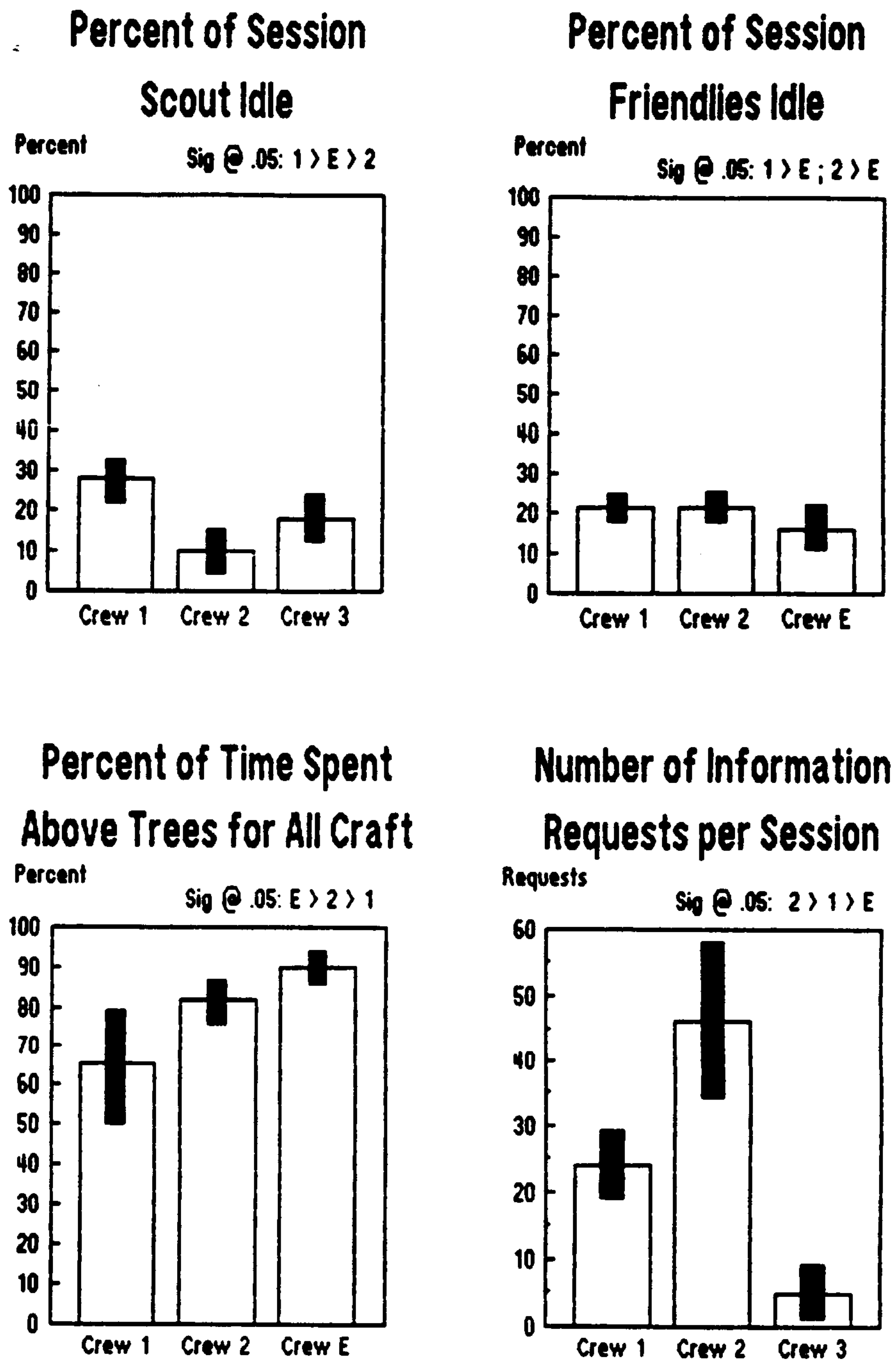

Crew comparisons of craft idleness, time above trees, and information requests

Figure 7

was superior to both the one-person crews in terms of minimizing scout idleness, it is surprising that Crew E was superior to the other two crews in terms of minimizing friendly craft idleness. This fact is peculiar given that Crew E had to time share control of the friendly craft with manual control of the scout, whereas Crew 2 had one crew member dedicated to friendly craft control, and Crew 1 used autopilot control of the scout and was thus presumably distracted from friendly craft control less often than was Crew E. Perhaps one reason for Crew E's superiority in this regard is the drastic difference in the number of information requests that were made for friendly craft status. This information consisted of resource levels (fuel, missiles, and weight capacity remaining) and craft altitude. While this major difference between the crews will be discussed in greater detail later, it appears to have been the case that Crew E was able to make decisions concerning friendly craft activity without explicitly consulting this information, resulting in more speedy decisions and thus, less friendly craft idleness.

The final graphs in Figure 8 suggest that, although Crew E was less able to keep the scout from idleness than was Crew 2, Crew E searched the world with the scout more completely than Crew 2. These facts would seem to suggest that the search paths generated for the scout by Crew E were superior in terms of the amount of unsearched area covered

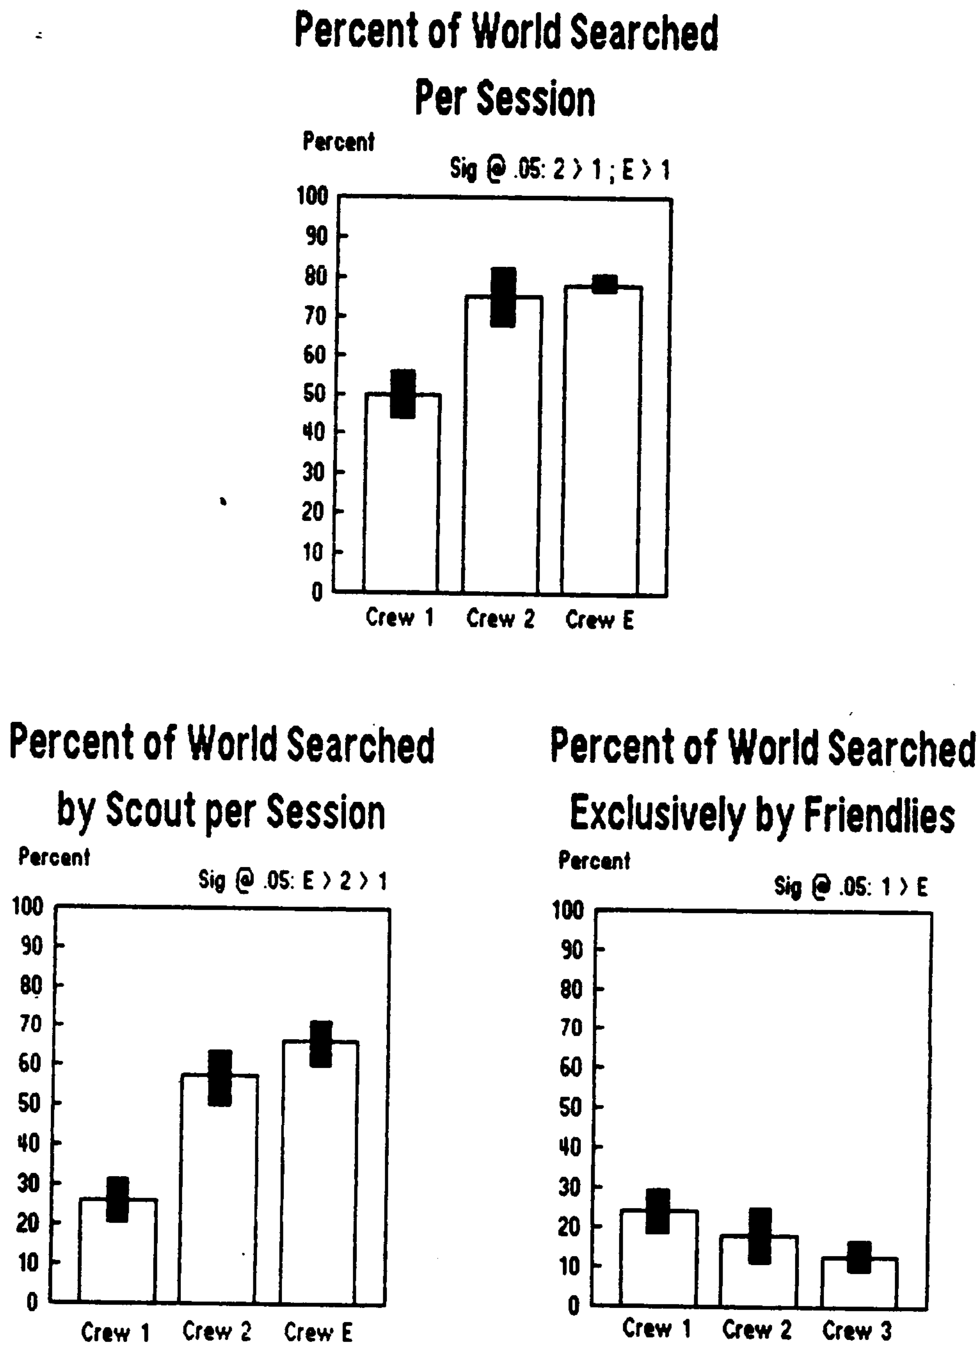

Crew comparison of searching effectiveness

Figure 8

per unit of time to the paths generated by Crew 2. Crew 1 was inferior to the other two crews in terms of both area searched in total and area searched by the scout. The fact that this crew covered more area exclusively with the friendly craft than the other two crews is probably due to the fact that this crew searched less area with the scout.

Summary and Hypotheses

It is clear that this analysis of performance is not equal to the task of clearly identifying all the strategic and competency related differences between the three crews. Nevertheless, some tentative generalizations can be made; they are only tentative since they will subsequently be evaluated in the modeling that follows.

It is clear that Crew 1 was inferior to the other crews on most performance measures. The primary reason appears to be an inability to effectively search the world with the scout. By using autopilot rather than manual control, this crew was able to maintain effective friendly craft control, but suffered in the long run since fewer cargo and enemy craft were discovered, and therefore the friendly craft could not be used to greatest benefit. Some of the scout's searching ineffectiveness was surely due to the increased fuel usage and slower speed under autopilot control, but it was also due to the fact that Crew 1 used the scout to load significantly more cargo than did the other one-person crew.

This reflects an insensitivity to the major functional difference between the scout and the friendly craft (search radius). On the other hand, Crew 1's behavior can also be viewed as a method for simplifying the task by only operating within a restricted section of the world. A second possible strategy is suggested by the fact that this crew never queued action scripts for the friendly craft, possibly indicating reactive behavior or behavior with a short planning horizon. This hypothesis is particularly difficult to evaluate since a lack of queued action scripts does not eliminate the possibility that future scripts were planned but not implemented via the text editor until they were needed.

Since Crews 2 and E were similar in terms of most of the overall performance measures, the major fact to be explained is how Crew E, as a one-person crew, attained performance equal to that of a two-person crew. It has been noted that although Crew E was less able to keep the scout in motion, he covered more unsearched area with the scout than did Crew 2. This fact suggests that Crew E's scout search paths were somehow more efficient. In addition, although Crew E had to time-share scout control with friendly craft control, he was able to minimize friendly craft idleness better than could Crew 2. Perhaps this fact is due to a more rapid decision-making process to determine friendly craft activity, as Crew E's lack of information requests would suggest. Finally, Crew E was seen to be more effective in bringing the cargo that were discovered to home base than was Crew 2. Perhaps this difference was due to a degree of coordination between the scout (which discovered most of the cargo) and the friendly craft (which carried most of the cargo home).

Clearly, the performance measures used above do not entail any of these hypotheses concerning crew differences. The generative modeling approach discussed in the following chapters was designed to provide a method by which these issues can be addressed more precisely.

CHAPTER IV

MODELING APPROACH

> You would be surprised how hard
> it often is to translate an action
> into thought. - Karl Kraus

Introduction

In this chapter the approach for modeling the three crews performing the laboratory task will be discussed. First, though, a brief statement concerning the focus of this investigation is made. This is done to restrict the scope of the problem to manageable size, to help define what other psychological research is relevant to the problem, and to identify candidate approaches for modeling crew performance. Then, other empirical and theoretical results concerning human information processing and skilled human performance are discussed that are relevant to the problem as defined. Then, an array of candidate modeling approaches are discussed and their strengths and weaknesses are identified with respect to demands of this modeling task. Finally, a general modeling framework for describing human perception, cognition, and action is developed. This framework will be made more specific and applied to crews performing the experimental task in the following chapter.

The Focus of Inquiry

Perhaps it is best to preface the review of existing psychological evidence by delineating exactly which psychological phenomena are of central importance to this investigation. This step is required due to the fact that the complexity of the task serves to raise a bewildering array of psychological issues. First, is should be clear from the data selection process (aggregating over the final sessions) that learning (whether perceptual, cognitive, or motor) will not be dealt with here. In addition, very "low level" or task-independent perceptual and motor processes are also excluded from consideration, as these phenomena are better studied in more tightly controlled experiments and more highly constrained tasks than the one used here.

The primary phenomenon of interest is how the three crews used the resources available (the four friendly craft and the scout) to accomplish the single goal of scoring points. Of lesser interest is how the crews "managed" their physical interaction with the interface in order to implement the friendly craft control activity, although, for the purposes of generative modeling, this issue must be addressed to some extent. Of major interest, then, is how crews decomposed the single goal of scoring points into a set of intermediate goals or tasks, how they constrained and coordinated the many degrees of freedom available for craft control, and how particular craft actions were chosen at particular times (both current and planned actions). The analysis of performance in the previous chapter suggested that crews may have differed along all these task dimensions.

It must be kept in mind that although the task of craft management is quite complex, there was not much time available for crews to perform their planning and decision-making tasks. This fact becomes clear when the number of action commands that were entered for the friendly craft via the text editor is taken into consideration. Crew 1 entered an average of 44.7 commands per session, Crew 2 entered an average of 67.4 commands per session, and Crew E averaged 51.7 friendly control commands per session. Taking 50 commands as an average, this means that a decision to change friendly craft activity was implemented, on average, every 36 seconds. This figure does not include decisions pertaining to control of the scout craft. In addition, the time interval between command implementations was not spent idly pondering information displays since many physical interactions with the interface (editing, maintaining scout control, e.g.) also had to be performed.

The picture of required control behavior that emerges is one of decision-making and planning that was rapid and, at least in the case of Crew E, performed simultaneously or rapidly time-shared with the scout control task that required (at least) perceptual-motor performance. The other two crews had considerably more time to perform the friendly craft control task (due either to the use of autopilot scout control or a second crew member) but neither of these crews were able to exceed the performance of Crew E. Thus, psychological mechanisms that are capable of producing effective decision-making and planning in the face of considerable task complexity and time pressure appear to be required to explain the behavior of skilled crews in this task.

Psychological Evidence

Due in part to the desire to produce "intelligent" machines, investigations aimed at understanding the differences between novice and expert behavior in a variety of cognitive tasks have become more numerous since the onset of the discipline of artificial intelligence. Those studies that are most interesting from the perspective of this work are those that require the human to solve complex information processing tasks rapidly. While some findings will be reviewed in this section, additional evidence will also be discussed in the following section where alternative modeling approaches are identified and evaluated.

Chase and Simon's (1973) studies of expert and novice chess players have suggested that the major difference between the two is in the amount and organization of prestored knowledge rather than in general information processes. Experts and novices were seen to be quite similar in terms of search depth, for example, but experts were found to be highly superior in terms of their ability to recognize chess configurations as similar and meaningful. Summarizing his own studies of chess experts, DeGroot has concluded that:

> The gist of the argument is that a chess position, and a fortiori an entire game are <u>typical</u> to the master. A chess position is easily recognized as one belonging to a certain class, that can be handled in a certain way. (1965)

Note the similarity of DeGroot's statement above with S. Dreyfus' characterization of skilled decision-making in complex managerial tasks:

> The significant pattern pervading the skill acquisition process, as we have described it, is the progression from abstract, rational under-standing in terms of isolated elements and rules relating them, to immediate situation recognition and response based on holistic similarity to prior concrete experiences. (1984)

Thus, at least in these two complex task domains, it appears as if highly skilled decision-making takes on a recognition and response flavor where the situations that are recognized are very complex, yet very quickly assessed.

While these two examples shed little light on the nature of the processes underlying situation recognition and response information processing, they do tend to suggest that such processing is more similar to perceptually oriented processing than the types of "higher level" cognitive information processing that are typically studied by cognitive scientists. Perceptual processing seems not to be restricted to simple visual discrimination and identification, but rather seems to be present in wide variety of tasks in addition to the chess and managerial decision-making tasks discussed above. The complex task of playing Go has been modeled as a perceptual task (Reitman, Nado, and Wilcox, 1978), and Smolensky (1986) has even described the task of intuiting answers to physics problems as perceptual in nature. The use of Chernoff faces for the display of multidimensional information (Jacob, 1978) is an example where the format of information presentation has been deliberately designed to promote a perceptual mode of processing instead of a serial and deliberate integration of the multiple dimensions.

Perceptual solutions to cognitive tasks tend to be characterized by high speed, parallel processing of spatial information (Shiffrin and Schneider, 1977; Treisman and Gelade, 1980), and a radical reduction in the number of decision alternatives that must be considered by "higher level" processing (Chase and Simon, op. cit.). It may be the case that, as Keil (1984) has commented, the human cognitive system may naturally gravitate toward this mode of processing and will, with experience, dispense with general purpose multi-stage processing routines and opt for reliance upon task-specific prestored knowledge of the type hypothesized by Chase and Simon. Anderson (1983) has gone so far as to suggest that the natural mode of information processing in the brain is highly parallel and pattern oriented.

While the previous characterization of skilled crew performance in the complex control task used in this research is possibly consistent with this perceptual style of information processing, it is quite clear that a perceptual approach is not entailed by the analysis of crew performance discussed previously. It would severely misrepresent the research reported here if a different set of empirical, albeit anecdotal, evidence was not reported that had a strong effect on the selection of a modeling approach. As mentioned previously, the present author served as an expert crew and trainer for two other crews in the third experiment. The additional evidence is his own experiences while serving the role of expert trainer and online advisor to these crews.

While watching over the shoulder of the two trainees, it was no more difficult to recommend desired control activity to the trainees than it was to determine control activity while performing "in the loop." What was exceedingly difficult was to explain the rationale behind those recommendations. What was even more difficult was to attempt to create a corpus of verbally represented procedural knowledge concerning a overall strategy for performing the task. When encountering any specific displayed situation within the context of the control task, decision-making was found to be quite easy, but in the absence of perceptual input, the task seemed foreign and it was difficult to make use of whatever was learned that accounted for highly skilled performance.

To make this frustrating experience more concrete, the reader should consider what his response would be if asked to write down a set of procedures for tying his shoelaces. The only way that many people report being able to create such procedures is to physically perform the task, and observe their own behavior. What is particularly relevant is not that the knowledge is somewhat inaccessible, but rather that the knowledge becomes immediately and almost effortlessly deployed when asked to deal with the task in a concrete, rather than abstract format. A more relevant and perhaps more tenuous observation is that this phenomenon does not appear to be restricted to the domain of perceptual-motor, or "low-level" tasks. For example, the reader may also consider his response to the demand to write down everything he knows about his own domain of expertise. This humbling and effortful task should be contrasted with the much more satisfying and easily performed task of solving a particular, concrete problem in that domain. Even in this case, what seems to be produced is not a stream of expert "knowledge", but a problem solution that only implicitly specifies what the expert must have "known" to produce the solution. The point of this exercise is to suggest that much of what accounts for expertise seems to be intimately keyed to the inputs of the information processing mechanisms, and cannot be neatly made distinct from perceptual input.

The common explanation for the phenomena of inaccessibility of expert knowledge is that that knowledge has become somehow "compiled" or automatized (cf. Ericsson and Simon, 1984; and Anderson, 1976, e.g.). Under this view, the phenomenon that accounts for skill is that intermediate steps in the "computation" are carried out without being independently evaluated or "interpreted", and thus, "[the] automation of performance is therefore quite analogous to executing a computer algorithm in compiled instead of interpreted mode" (Ericsson and Simon, 1984). What is very important to note about such theorizing is that all of the intermediate steps in the original procedure are preserved in the automatized procedure, although some of them are no longer capable of being verbalized. That is, skilled routines are simply unskilled routines that are speeded up due to the fact that their inputs and outputs no longer require use of short term memory (Ericsson and Simon, 1984). The information processes underlying skilled performance are thus no different in kind from unskilled processes, other than for the fact that they have been stipulated not to require use of short term memory.

What these explanations typically fail to explain, though, is how expert knowledge becomes so effortlessly tapped when the expert is presented with a concrete problem (or else he would not be considered an expert). The explanation offered here, and one that seems to be consistent with the findings discussed above, is that much of this "knowledge" is often comprised of perceptually oriented routines that are effective only when working within their "design specifications," i.e. when processing input information. These perceptually-oriented routines are different in kind from unskilled processes which do not make use of a perceptual processing mode. The inaccessibility of skilled processes is not explained by stipulating that they do not use short term memory as under the automatization theory, but, rather, it is explained by the fact that a shift to a perceptual mode of processing will exhibit those features that are typical of perceptual processing, such as inaccessibility, high speed, and parallel processing. While such a shift may indeed result in a lessened use of short term memory, this fact need not be seen as a defining feature of expert processing, but rather as a feature that is exhibited due to a shift in processing mode. To summarize, the offered view of the transition to skilled processing is to be characterized as a shift to a different kind (perceptual) of processing, rather than as the same processing occurring in unskilled processing that is simply speeded up and stipulated not to require memory resources. Of course, this view can be no more precise than the characterization of perceptually-oriented processing, and this is one of the fundamental goals of the modeling exercise.

It should perhaps not come as a surprise that much of what accounts for high skill might involve becoming attuned to the perceptually available specifics of the situations that are encountered in a task. It is a widely held principle in engineering psychology that, with experience, human behavior becomes a reflection of the task environment (Rasmussen, 1983). In a wider domain it has been suggested by the philosopher Jaspers that people in general become "their situations personified," (Murphy, 1978), and, in an even wider domain, evolutionary theory suggests that all organisms become intimately attuned to the demands of their ecological niche. If this line of reasoning is close to being correct with respect to modeling skilled performance, some rather severe restrictions must be placed on the models that are held to be candidates for descriptions of human behavior in complex systems. Some such models are reviewed in the following section.

Candidate Modeling Approaches

One class of models that has been proposed to describe human behavior in complex supervisory control systems is based on normative constructs from decision theory, optimal control theory, and operations research. Normative models of the skilled human controller based on utility theory and dynamic programming have been developed by Sheridan (1976 and 1966, respectively). An explicitly stated assumption of Sheridan's (1976) model is that "the human operator's (and/or computer's) purpose is to maximize an explicit or implicit utility function." A model based on queuing theory has been developed by Rouse (1977) to determine dynamic task allocation in multitask environments. The PROCRU model developed by Baron, Zacharias, Muraldiharan and Lancraft (1980) is an adaptation of the optimal control model that has been tailored to include procedural information.

In addition to being normative, all these models share another important feature in common: they abstract away the contextual features of the decision-making task and focus on abstract, formal features of the decision problem. This is in keeping with the general approach of the decision sciences to find solution methods to problems across the many domains in which those problems may be "embedded". As such, these methods categorize problems by their methods of solution, and necessarily dispense with the particulars of the decision task. Thus, this approach tries to find the same solution methods to problems in various contexts, where the solution process that is selected is based on an abstract and formal characterization of the problem.

If the remarks of the previous section concerning skilled human decision processes have any validity, the features of human decision-making might suggest a decomposition of problem solution methods along dimensions nearly orthogonal to the dimensions respected by these normative models. Rather than coming equipped with an array of various decision-making methods that are variously selected by abstract features of the problem task, and applied independently of context, humans may employ an array of approaches that are selected primarily by context, and are applied nearly independently of whatever formal structure the decision task may have. Thus, for example, while the decision sciences may employ different solution techniques for choice, inference, or planning tasks, due to differences in formal problem structure, skilled human cognition may gravitate toward the same solution technique (situation recognition and response) given that this technique is enabled by the specifics of the context in which the problem occurs. On the other hand, the decision theoretic approach might adopt a linear programming solution technique, for example, for a task with the appropriate structure, regardless of the context in which it appears. Human cognition, though, may not respect this abstract invariance in problem structure and adopt different solution techniques depending on what types of information processing are most readily supported by the context in which the problem occurs.

A vast array of psychological findings attest to the context-specificity of human information processing. Context-specific processing is characterized by sensitivity to the surface (or perceptually available) aspects of an information processing problem. This type of processing is identified when variance in behavior is exhibited due to altering way in which the same (when formally characterized) problem is presented to the human. Many studies have shown that a set of problems with identical formal structure in some abstract problem representation are attacked quite differently by humans depending on the context in which the problem is couched. In studies of reasoning, Wason and Johnson-Laird (1972) and Johnson-Laird (1975) have demonstrated that the inferences humans draw, and the validity of these inferences, when presented with a task of invariant syllogistic structure are strongly dependent upon the context in which the task is presented. Studies of inference and decision-making by Kahneman and Tversky (1979) and Tversky and Kahneman (1981) describe considerable sensitivity to contextual information over tasks that are identically represented in terms of probability theory. In their study of problem solving, Kotovsky, Hayes, and Simon (1985) found reliable performance differences of a factor of sixteen to one in the Tower of Hanoi problem when the problem context was altered, even though the problems had identical structure in the "problem space plus operators" problem representation developed by Newell and Simon (1972).

These facts seem to call into question the psychological validity of models such as those based on normative constructs from decision theory, queuing theory, dynamic programming, or optimal control theory. By abstracting away contextual information, none of these approaches may be suitable for adequately describing skilled human decision-making processes that are notoriously context-sensitive. The necessity of addressing the context-sensitive nature of human information processing could not be more important than in the discipline of human-machine systems design, since the problem of interface design is largely one of designing a (hopefully supportive) context for human information processing to occur.

A second class of models has been developed to explicitly account for the domain-specificity of human information processing. Expert systems and other symbol processing oriented models of cognition make use of the fact that expert cognition is strongly determined by the properties of the particular domain of expertise. As a result, these models are often termed "knowledge-based" systems. These models often encode human expertise in terms of verbally represented rules and facts concerning the problem domain. The generative solution technique is the application of logical deduction or induction. Examples of this approach used to describe cognition in complex dynamic systems include, for example, Chen (1985), Wesson (1977), and Rouse, Rouse, and Pellegrino (1980). The models constructed by Chen and Wesson are knowledge based systems that generate plans in the domain of flight planning. Wesson has suggested that his model could be used as an aid for the human air traffic controller. The model built by Rouse et. al. is a description of human problem solving in a dynamic system. Unlike the flight planning models, though, this model was demonstrated to generate behavior in approximate agreement with human behavior.

There are two features of these models that make them questionable from a psychological perspective, even assuming they can produce behaviorally valid results. The first is due to the fact that although these models are indeed highly domain dependent, they are still context-insensitive. These models do represent knowledge that is highly domain specific, but the representational form that is used for this knowledge is almost invariably verbal, and thus divorced from whatever perceptual processes the domain expert might use to encode and deploy this knowledge. If some expertise involves perceiving complex patterns within the environment, it is not at all certain that natural language has the descriptive resources to economically represent these patterns and support the processes that recognize them. These models assume that the knowledge accounting for expertise is not only verbally encodable (as much knowledge may be, as expert interviews may suggest), but that these verbal representations are sufficient for the task of capturing the expert's internal information processing.

It seems to be an almost unquestioned assumption in cognitive science research that the (verbal) representational system underlying communication is the same representational system underlying cognition. Given that the functional roles of communication and cognition are radically different (unless cognition is equated with "talking to oneself"), there is no a priori reason to expect that these two processes should share the same representational system. Do cats "think" by manipulating various internally represented meow sounds, and does a vocabulary of high pitched songs serve as the "language of thought" for the whale? (See Churchland, 1986) While some human cognition might be profitably be described as inner speech, the cognitive processes of the "participants" to the discussion still remain to be accounted for. In other words, it would appear that a theory equating cognition with inner speech requires the assumption of mechanisms that produce and interpret this speech. It seems clear that much of the psychologically significant content of such a theory would have to be embodied in assumptions about these mechanisms. But the operation of these mechanisms cannot once again be described as involving inner speech without the possibility of an infinite regress.

If, as the remarks of the previous section suggest, much of expert processing may be perceptually oriented and may make explicit and profitable use of the context in which the perceptual system operates, these verbal representations may not provide the appropriate representational language. It may be the case that, as Nisbett and Wilson (1977) suggest, that the verbal statements offered by domain experts in response to queries by knowledge engineers are based on a priori, causal theories that plausibly explain stimulus-response relationships. In the shoelace tying experiment for example, the verbal rules that the reader was asked to generate can be seen as describing certain input-output invariances that were noticed while observing shoelace tying behavior, but there is no guarantee that the reader was actually reporting rules that played a generative role in the production of that behavior.

In addition to representational form, the generative solution mechanism employed in these procedural models is also context-insensitive. While, at a high level of description, these models can be described as using domain dependent solution strategies, at the level of mechanism all these models share an identical generative technique: the application of logical inference. One explanation for the ubiquity of the logical model of cognition is evidenced in Churchland's (1986) comment that:

> I suspect that the philosophical tradition of venerance for inference and the sentential

> (verbal-propositional) attitude has generated a kind of fetishism with respect to logic as the model for inner processes.

There is simply no evidence that the logically based model which adequately captures some human reasoning at the <u>personal</u> level (the verbally reportable manipulation of syllogisms in philosophical argumentation, e.g.) is well suited to capturing expert cognitive processes that are not verbally reportable and occur at the <u>sub-personal</u>, or cognitive level.

The distinction between personal and sub-personal levels of cognition is meant to reflect a category distinction that sometimes appears to be blurred in cognitive theorizing. While a person may be usefully described as following a rule, (when cooking according to a recipe for the first time, or when attempting to troubleshoot a car's ignition system with Chilton's manual in hand, e.g.), or making a deduction, (by manipulating syllogisms in a freshman logic course, e.g.), the cognitive processing underlying this behavior cannot be generatively described as rule following or logical deduction without (explicitly or implicitly) invoking circular reasoning. It appears somewhat idle (from a generative perspective) to explain the human's ability to follow a rule by contending that his mind/brain is rule following, or to explain human logical deduction by postulating a logical mind. These arguments contain little more meaningful content than, for example, the explanation that a calculator can be described as following the axioms of arithmetic because its arithmetic chip is also following those axioms. Just as certain properties (obeys arithmetic axioms) apply to calculators and a different set of properties apply to chips (is wired in a certain way) by virtue of which the calculator has this property, different sets of properties appear to be required to describe personal-level behavior and cognitive-level behavior in order for explanations of the former in terms of the latter to be productive and non-circular. For additional arguments against the view that cognitive processes are best viewed as logical manipulations of verbal representations see Coulter (1983), Winograd and Flores (1986), Churchland (1986), Hunter (1973), Runeson (1977), Stich (1983), Rumelhart and McClelland (1986), and Hinton and Anderson (1981).

It thus appears that the most cautious approach demands that the use of logical-verbal models should be restricted to those cases where there is evidence that logical inference on verbally encoded information is, in fact, being employed by the human to perform a task. In these cases, it may prove to be a productive research strategy to use logic to provide a descriptive account of behavior and to ignore the lower level cognitive mechanisms that occur, and by virtue of which, the human can exhibit this logical behavior. Human problem solving and plan construction in novel situations may, for example, be phenomena that meet these criteria. In other cases, though, where there is little evidence to suggest explicit logical inferences are operative, other characterizations of cognition may be necessary.

Modeling Approach

The modeling approach is based on the view, suggested above, that the boundary between perception and cognition is a shifting one. Answers to problems that once had to be "worked out" can, with experience, be "seen"; a change not unlike the shift that occurs when learning to read a second language. The perceptual mechanisms responsible for such a change are assumed to be differentiated by the information to which they are attuned and the answers they supply, rather than by the formal characterization of the problem (inference, planning, maximizing utility, e.g.) that they can be described as solving. That is, there is a unitary solution mechanism (seeing answers) that is not only context-sensitive, it is context determined, and somewhat divorced from whatever formal structure the task may have when defined from the perspective of the decision sciences. The formal and reward structures of the task will be relevant in determining which specific perceptual mechanisms are developed, so that the operation of these mechanisms can contribute to behavior that satisfies these formal constraints. While these task-specific perceptual mechanisms have most surely evolved under the pressures of reward, punishment, or some similar shaping forces that are sensitive to task goals and result in expert performance, an appropriate characterization of the operation of these mechanisms is essentially ahistorical and non-teleological. In other words, a description of the real-time, physical operation of these mechanisms need not refer to the past events that were responsible for their existence, or the goals to which their operation contributes. An automobile, for example, can be described within the domain of physics, without knowledge of the marketing research that led to its particular design, or knowledge of the institutional goals that the production of the automobile was intended to serve.

Under this view, any successes of decision theoretic modeling of expert performance in perceptually rich domains are to be explained by suggesting that perceptual mechanisms have been evolved to supply the optimal solutions, although no optimization operations are being performed during the time period in which the answer is supplied. The descriptive successes of rule-based and other symbol processing models for describing highly skilled performance are to be explained by contending that the rules and logical models are higher level descriptions of perceptually-based behavior, and for some purposes perhaps even redundant descriptions.

This is a radically strong statement of the ideas on which the modeling approach is based. Even if these ideas are on the right track, perhaps the above statements are only even close to being strictly true in extremely highly skilled behavior that occurs in perceptually rich domains. Nevertheless, the modeling approach has been stated in this fashion because, while the ideas are quite simple, it is a direct implication of these ideas that any task-specific model based on these ideas must necessarily reflect all of the task's complexities and idiosyncracies, and thus, will not communicate the simplicity of the approach. This is especially true of the generative model described in the next chapter.

Given this perceptual characterization of skilled performance, two major questions need to be addressed before the framework can be applied to this or any similarly complex task. First, which perceptual processors will exist, given the task requirements and information available to the human? Second, what is the mechanism by which these processors operate? The answer to the first question will determine which input-output mappings are produced, while the answer to the second will determine the nature of the processes that actually generate these mappings. It is assumed that the nature of the computations producing the mappings (at some level lower than a purely functional level) is independent of the particular input-output mapping that is being generated.

It is clear that the task-specific perceptual mechanisms that occur will be determined by the requirements of the task that is being performed. At the most general level, the task is action. While this may seem to be a trivial point, it is clear that if it is contended that these mechanisms help account for skilled performance, rather than simply skilled perception, the role of these mechanisms is more than to simply register, or provide a faithful internal representation of, the external environment. Rather, the role of the perceptual mechanisms is to provide information upon which the selection of action can be efficiently performed. An appropriate model of human information processing requires knowledge of what action related, task-specific information these mechanisms use and produce, as Neisser suggests below.

> If we do not have a good account of the information that perceivers are actually using, our hypothetical models of their "information processing" are almost sure to be wrong. If we <u>do</u> have such an account, however, such models may turn out to be almost unnecessary. (Neisser, 1987)

Neisser's comment is representative of the ecological approach to perception and cognition. The ecological approach, stimulated by Gibson's work in perception, seeks to find much more information in the environment than is usually assumed by the more traditional information processing approach (Gibson, 1966). Under this view, perceptual systems tuned to complex informational invariants can serve to reduce much of the computational burden for the production of complex behavior from central processing. Although the ecological view has spawned an impressive amount of research activity, it is still far from uncontroversial (for criticisms, see Fodor, 1983, and Pylyshyn, 1981).

For the purposes of this work, employing an ecologically oriented approach to modeling human behavior implies only the following commitments: first, the information provided by perceptual mechanisms is highly action-oriented; and second, these mechanisms operate by being attuned to (possibly complex) invariants present in the environmental situation. On the other hand, the present approach may part company with ecological views that assume that the pick-up of information is direct, or totally unmediated by any computation or information processing.

Action-oriented information, or "affordances" (Gibson, 1966), are properties of the interaction between an environment and an organism's capability for action. For example, a chair affords sitting, an apple affords eating, and a cigarette affords smoking. While clear enough, these examples are somewhat misleading because they are suggestive of an environment that is already neatly differentiated into these objects, and the fact that they afford different types of actions is not essential to the differentiation.

A more generative conception of the affordance relation (and the one adopted here) is that an affordance is not a benign element that exists after the environment is differentiated, rather, it is a partial source of the differentiation. That is, the environment is differentiated by the degree to which it variously affords actions of different type. A chess player, recalling DeGroot's characterization, decomposes the chess game into situations that "can be handled in a certain way" thus reflecting an action-oriented differentiation of his environment. This idea is consistent with Rosch's influential ecologically oriented work on human categorization where one feature of "basic categories" is that they are comprised of objects that support human actions of specified types (Rosch, 1975). In addition, an affordance need not be an all or none property, as situations may differ in the degree to which they afford productive action (fishing, mining, surfing, e.g.).

The idea that environments become differentiated in this way is especially productive in situations where a decomposition of the environment along solely physically salient dimensions (color, shape, etc.) is not especially efficient with respect to the task of action selection. In some cases, a purely physically oriented decomposition will under-differentiate the environment. In these cases, ignorance concerning possible action-oriented differentiations can make behavior appear very complex or even random. Obtaining an appropriate differentiation can sometimes allow such complex behavior to be described in simple terms, as will be seen in the model of human path planning described in the following chapter.

Sometimes, on the other hand, a purely physically based environmental decomposition will over-differentiate the environment with respect to the human's capability for action. In these cases, affordances will determine how the perceptually distinct objects in the environment become categorized. That is, one way that a set of objects can be partitioned into categories is by classing them according to the actions they afford. The same set of objects can be categorized in different ways depending upon the action affordance that is employed to generate the decomposition. Knowledge of actions is required to identify the categories that are formed. For example, the category (children, jewelry, legal documents, and cameras) is almost non-sensical until it is recognized that the members afford actions of the same type; in this case, they are "things to take from one's home during a fire" (Barsalou, 1985). Thus, perceptual processors that are individuated by their sensitivity to affordance-related information are assumed. The model discussed in the following chapter is based on the idea that the crews differentiated and categorized their environment (the simulated world) by the degree to which situations present in this world afforded actions of various types (searching, loading cargo, attacking enemies, e.g.).

The question of the mechanisms by which these processors are assumed to operate must now be discussed. As mentioned earlier, it is assumed that these mechanisms operate by processing spatially distributed information rapidly and in parallel. In addition, mechanisms based on logical operations on verbal representations are rejected, although the operation of these mechanisms may be consistent with a such a high level description.

The parallel distributed processing (PDP) approach (e.g., Rumelhart and McClelland, 1986; Hinton and Anderson, 1981) for describing cognition employs perceptual and cognitive mechanisms that appear to be consistent the processing requirements specified above. This approach assumes that input information is processed in parallel by a set of spatially distributed processors that are interconnected by simple communication mechanisms. The operation of these mechanisms are typically described with a combination of matrix algebra and the use of simple non-linear thresholding functions. While this approach has achieved some successes for describing perceptual abstraction and categorization, its utility for describing the processes that account for complex human behavior is still in question (Fodor and Pylyshyn, 1988; Norman, 1986).

Certain of the perceptual mechanisms employed in the following model are consistent with these PDP mechanisms in that transformation of vectors via simple matrix operations are used. Other mechanisms, on the other hand, are implemented via simple logical operations along with matrix algebraic methods. As logically based descriptions can provide a higher level approximate description of the behavior of PDP mechanisms, the use of logic for implementing the model is only done for programming ease only. The entire model structure is still based on the assumption that these mechanisms operate by processing the spatially distributed information in parallel without the manipulation of verbally based representations. Consistency with this assumption is the major motivation of the modeling approach, rather than any desire to maintain similarity with the current assumptions underlying PDP modeling.

A final point should be made concerning the relationship of PDP modeling with the present approach. Recall that a primary motivation of the present approach is that skilled human cognition is extremely sensitive to the context in which an information processing problem is presented. PDP models are also completely context sensitive, as they operate by processing concrete, primitive perceptual features, and may require completely different input-output mappings when the format of information presentation is altered. This feature makes PDP models consistent with the present modeling approach. It has also been suggested that the context-sensitivity of human information processing indicates that models based on abstract, formal problem representations do not provide adequate descriptive accounts of cognition. This is a particularly distressing fact for the human-machine systems researcher because, if true, invariances in problem structure that are respected by the formal decision sciences may not be similarly respected by human cognition.

What is suggested then, is that the search for invariances in cognition that will have predictive value for the system designer may need to change from the level of performance mechanisms to the level of the processes by which the context-specific mechanisms are produced and tailored to the interface design and task demands. That is, the search for predictively valid invariances might need to shift to the level of learning, otherwise the task-specific perceptual and cognitive structures that are evolved cannot easily be determined. The performance-oriented parameters of the structures can only be explored after the task-specific structures have been identified. That the phenomena of learning, then, occupies center stage in the PDP approach to cognition may be a reflection of the validity of this observation.

CHAPTER V

MODEL DESCRIPTION

Introduction

In this chapter, the generative model used to describe the behavior of subjects in the laboratory task will be described. The model has three main components: perceptual, action selection, and action implementation mechanisms. First, a qualitative description of the perceptual processing employed in the model is given in terms of the framework developed in the previous chapter. Then, the structural organization of the model is discussed which includes the perceptual, action selection, and action implementation mechanisms. Next, the operation of each of the perceptual, action selection, and action implementation mechanisms is described in detail. The methods by which the model was used to describe individual crews are described in the following chapter. The results of crew modeling are discussed in Chapter VII.

Framework for Perceptual Processing

In keeping with the remarks of the previous chapter, it is assumed that much of the information processing accounting for highly skilled performance in this task is performed by task-specific perceptual processing mechanisms. The particular perceptual mechanisms that are developed are determined by the task requirements and the interface displays that provide the context in which the task is performed. More specifically, the perceptual mechanisms are assumed to be attuned to the features of the displayed world that are highly relevant for the purpose of action selection. These mechanisms are assumed to process spatially distributed information rapidly and in parallel.

The environment is assumed to become differentiated and categorized by the degree to which it affords actions of different types. Differentiation occurs when an originally isotropic description of the environment becomes enriched due to an inclusion of information concerning how different sections (sometimes "objects") in the environment afford a specified action to various degrees. Sighting cargo and enemy craft, for example, is an action that is afforded to different degrees by different locations in the simulated world due to the fact that object densities varied with forest density. Locomotion of the scout is a similar such action, as ease of locomotion is determined by the forest density in the immediate location of the scout craft. The complex action of searching for objects includes both sighting and locomoting as simultaneously performed subactions, and the differentiation of the environment according to this complex action can be constructed from the two differentiations provided by its component actions. Therefore, a complex differentiation of the world is provided by search affordances. This differentiation is not based solely on the physical attributes of the world; rather, it is generated by considering how the world's physical attributes combine with the functional capabilities of the scout to produce regions of high and low value. The resulting differentiation does not decompose the world into distinct objects, but rather, it induces a continuously graded affordance structure that includes "hills" and "ridges" of high degrees of search affordance, and "valleys" and "holes" of low degrees of search affordance.

Categorization occurs when the environmental decomposition based solely upon physical dimensions (e.g., by color and shape) is reduced in dimensionality by noting that various items in this decomposition all afford the same action. For example, the model exploits a categorization of both enemy helicopters and enemy tanks into the single category of mobile enemies due to the fact that both these objects are acted upon in an identical way (with a given friendly craft action command) even though the physically generated categorization based on color groups these items into separate classes (helicopters appear as red circles whereas tanks appear as orange circles). The model makes use of a different category for fixed ground enemy craft, as these enemies are attacked with a different command than is used for the mobile craft. All categorization that operates by abstracting away the specific numbering of these enemy craft operates in a similar way, as enemy craft are attacked in the same way regardless of their numeric identifier.

Sometimes, the object categorization based solely upon physical dimensions is nearly identical to an affordance oriented categorization. For example, cargo are the only gray objects on the display, and each and every cargo (to close approximation) is acted upon in the same way (with the same friendly craft cargo loading command). This command directs the friendly craft to travel to the cargo location and load the cargo. The only exception to this categorization arises when cargo appear within the loading distance (0.125 miles) of a friendly craft. In this case, the friendly need not be given a command to travel to the cargo location, as this constraint is already satisfied. Therefore, a different action command is executed to load these craft. This suggests that there might be a corresponding perceptual mechanism that is sensitive to those cargo that are close to friendly craft, as they afford actions of a different type than cargo that require that friendly craft travel.

Note that a comparison of the purely physically based categories and differentiations with the affordance based categories and differentiations provides one way to analyze the compatibility characteristics of displays, and helps define which subtasks are variably mapped and which are consistently mapped (Shiffrin and Schneider, 1977). This latter distinction can be used to predict which tasks will become automatized, and which resist automatization. This distinction will be discussed more precisely in following chapters.

Similar reasoning is used to identify the entire set of perceptual mechanisms and the hypothesized affordances to which they are sensitive. Perhaps most interesting are those affordances to which the model is sensitive that result in coordination among the friendly craft, as coordination is an example of one of the behaviors that seems to be suggestive of "higher level", central, or more flexible cognitive operations. The general mechanism by which coordination is produced in the model is by allowing action affordances for any one friendly craft to be partially composed of information relating to the activities of the other friendly craft. In short, part of the environment of each craft is comprised by the other four craft.

Examples of coordinated behavior among the scout and four friendly craft produced by the model include collision avoidance, making sure friendlies do not attempt the same action (e.g., loading the same cargo, searching the same region), and coordinating the space-time trajectories of the scout and other friendly craft so that a friendly craft is nearby to process any cargo and enemy craft sighted by the scout's radar. All these behaviors are produced in a similar fashion, namely, by allowing the affordances for one friendly craft to be determined by the current and predicted activities of the other craft. For example, in the scout-friendly craft coordination task, areas that simultaneously are capable of containing both the scout and friendly craft have a high degree of search affordance for the friendly craft. The search affordance distribution for a given craft, then, is partially constructed from the current and predicted location of the scout craft. A similar method is used to avoid collisions between the five craft, to coord-inate times at which the five craft return to home base, and to make sure that each object and search region is allocated to at most one friendly craft.

Another example of "higher level" behavior that is generated by sensitivity to environmental affordances in the model is planning the future activities of the five craft. For planning, the identical set of perceptual mechanisms that provide for current craft action selection is used. What is changed, though, is the information to which these mechanisms are applied. The model selects future activities for each of the craft by forward simulation of the world based on difference equations that specify world dynamics. These equations describe how the world state (the existence of cargo and enemy craft) and the friendly craft state (locations, fuel levels, missiles, cargo carried) change over time due to scout and friendly craft activities. Using these equations, the world is essentially advanced to a future time point, and the perceptual mechanisms operate by assessing the affordances that will exist in this future world. The relevant future time points are the end points of current or already planned activity. The psychological mechanism that is assumed to provide this function for the human crews is visual imagery. There is some evidence (and much more conjecture) that the same perceptual mechanisms that are responsible for abstracting information from the environment can be partially stimulated in a top-down mode to produce visual imagery (Neisser, 1976).

The selection of friendly craft is quite simple: each craft is commanded to take the action that has the highest affordance. As described in the previous chapter, an affordance is a relation of the interaction between the environment and an organism's capability for action. The way in which the environment can differentially support actions has been mentioned above, but the organism's side of the relation needs to be discussed. One important set of information to which friendly action decisions must be sensitive is each craft's resources (fuel, missiles, weight carrying capacity). Sensitivity to this information is nicely couched in ecological terms by noting that a craft's resource levels essentially determine the craft's capability for action. That is, the different actions that a craft is capable of executing at a particular time will be partially determined by the craft's resource levels. For example, whether or not a craft is capable of loading a cargo or attacking an enemy is affected by its weight capacity and number of missiles. The resource information, then, comprises the craft's capability for action, the environmental structure determines the degree to which each of these actions is provided by the environment, and the affordance for a given action is a function of the interaction between these two factors. High affordance levels arise when the craft is capable of performing an action made available by the environment, and low affordance levels are characteristic of a mismatch between a craft's capabilities and the environmental structure.

Thus, a sensitivity to the affordances in the simulated world is the unitary, context-specific mechanism by which the model performs the tasks of decision-making, coordination, and planning. Few of the formal features that would typically be exploited by operations research or artificial intelligence methods to differentiate these information processing problems are reflected in the model. On the other hand, the action selection components of the model can be interpreted within the framework of multi-attribute utility theory (MAUT) descriptive models of human decision-making (Raiffa, 1970). The similarity is due to the fact that the processes used here to integrate the affordance levels of sub-actions into a resultant affordance level for the total action can be interpreted as combining multiple utilities for the individual attributes of the action. The modeling approach adopted here for action selection, then, need not be considered at odds with MAUT modeling. Rather it can be viewed as a way of supplementing the MAUT approach by aiding in the identification of the relevant problem attributes via viewing the problem in affordance oriented terms, and by indicating how the problems of choice, planning and coordination can be approached within a single framework and in a manner consistent with the characterization of the information processing underlying skilled performance that was discussed in the previous chapter.

Structural Description

The model structure can be decomposed into three major components: the perceptual mechanisms, the selection mechanisms, and the action mechanisms. The perceptual mechanisms, defined qualitatively in the previous section, process the displayed information and produce affordance oriented information concerning the attractiveness of each of the actions available to the scout and each of the friendly craft. The selection mechanism accepts the affordance information for each of the craft and determines required present and future actions. The action mechanisms accept the required craft actions and are responsible for simulating the physical and visual operator activities at the interface that are required to implement these actions.

Figure 9 on the following page schematically depicts the model decomposition into perceptual, selection, and action processors. For the perceptual and selection processors, the diagram also indicates the individual mechanisms that make up each type. The action mechanisms are not decomposed in this diagram. A detailed description of the entire set of action mechanisms will be provided separately in a later section of this chapter. While each of the mechanisms in the diagram will be described in more detail in the following section, it is necessary to first describe them at a gross functional level to indicate the global organization of the model.

The two top perceptual mechanisms, the scout search affordance and scout locomotion affordance mechanisms provide the information upon which the scout search path selection mechanism determines search paths for the scout craft through the simulated world. The search affordance mechanism operates by creating a spatially distributed affordance mapping of the world, where the affordance of a world location is determined by the effectiveness of scout radar centered at that location for sighting possible cargo and enemy craft. The locomotion affordance mechanism creates a spatial mapping of the world, where the affordance

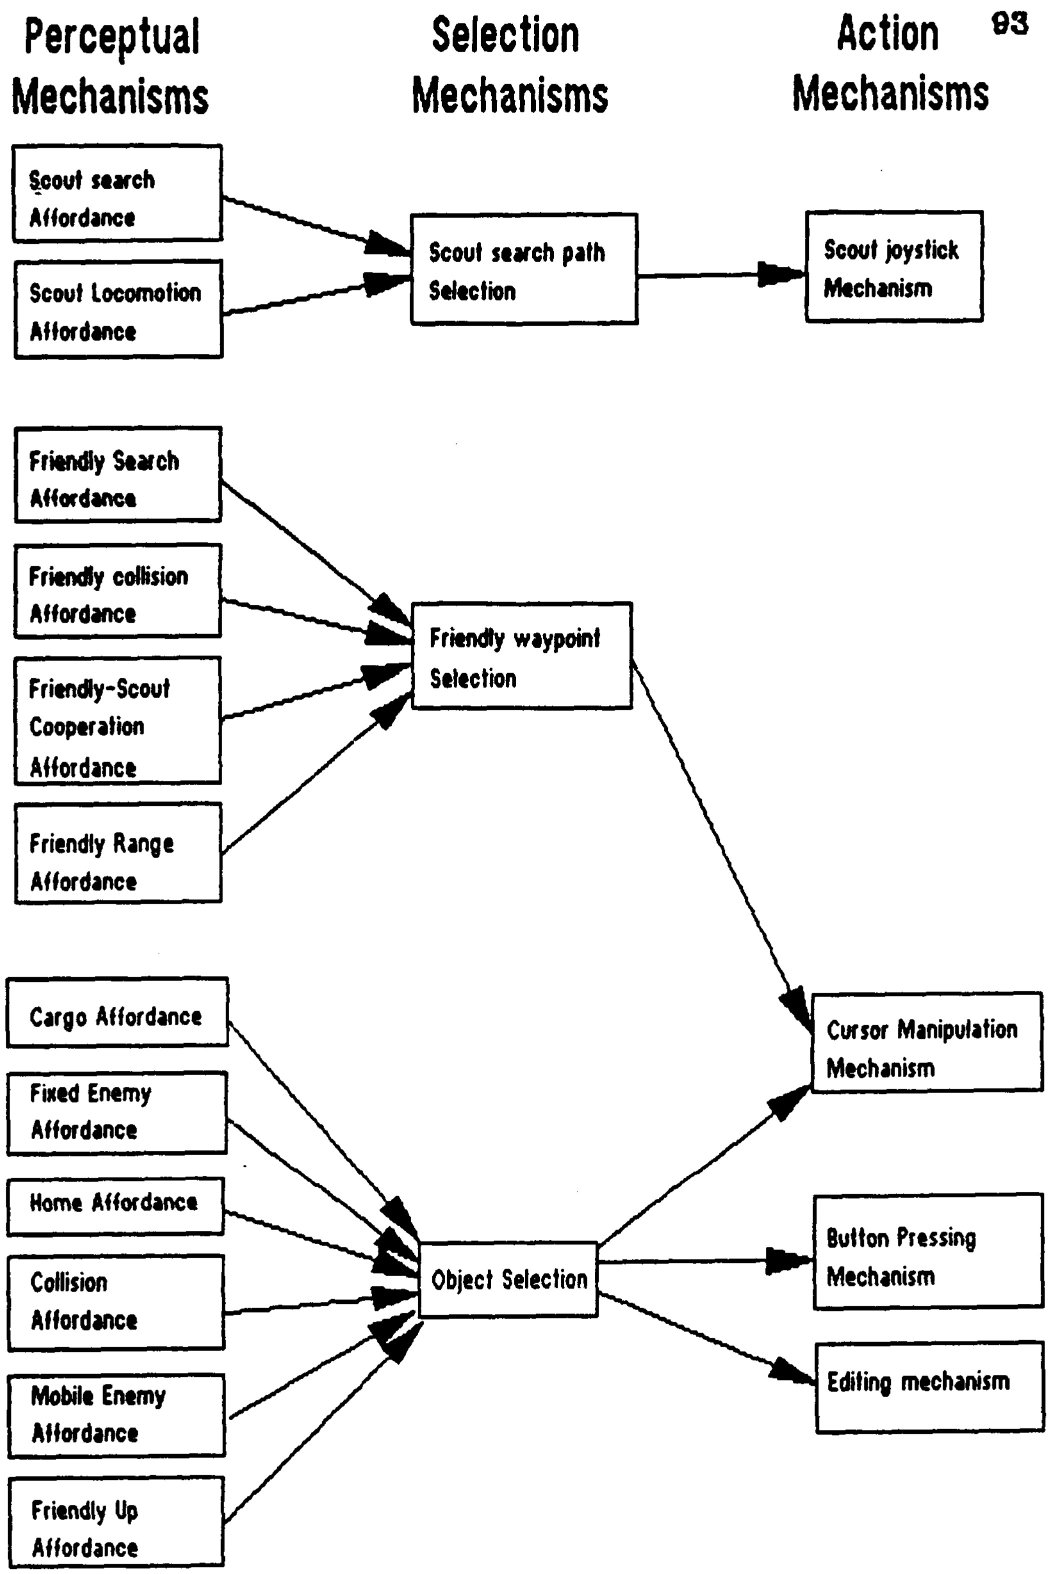

The organization of perceptual,
selection, and action mechanisms

Figure 9

of a world location is determined by the speed with which the scout can be flown through the trees at that location. The scout waypoint selection mechanism operates on the additive combination of these two affordance maps and selects a desired next waypoint for the scout craft. The primary output of this selection mechanism is to the scout action mechanism which implements the scout search plan by performing the control actions necessary to actually move the scout to the desired waypoint. In a complete model, this task would be performed by a control theoretic model capable of accepting a desired waypoint and producing the manual control activity necessary for reaching it. In the present model, though, this task is approximated by simply moving the scout in a series of very short steps (0.06 miles each). The method by which these steps are determined will be discussed in the description of the scout motion action mechanism in a later section.

The next four perceptual processors in Figure 9 are responsible for providing the affordance information upon which search waypoints for the four friendly craft are selected. The friendly search affordance processor produces a map similar to the scout search affordance map. The affordance level of a world location on the friendly search affordance map indicates the desirability of searching the area that would be covered by friendly craft radar along the linear path to the waypoint location. The affordance is a function of the linear path because the selection of a waypoint is actually the selection of the linear route along which the friendly craft should travel. The friendly collision affordance processor creates a spatially distributed affordance map that signals areas of potential collision with other friendly craft. The scout-friendly coordination processor creates a spatially distributed affordance map that indicates regions of high value due to the fact that they are near the current or predicted scout craft location. The friendly range affordance processor produces a spatial map of the world that indicates the areas in the world that are in range of the friendly craft; i.e. those to which the friendly craft may travel and still have enough fuel to either last till the end of the mission or return home to resupply. The friendly waypoint selection processor accepts the combination of these four affordance maps and selects a search waypoint for the craft.

The bottom six perceptual mechanisms produce affordance information upon which object (cargo, enemy craft, and home base) selection decisions are made. The cargo affordance indicates the degree to which a cargo at a specified location affords being loaded by a friendly at a specified location with a specified amount of resources. Considerations in this affordance calculation include: 1. the distance between the friendly and the cargo; 2. whether the friendly has enough fuel to load the cargo and take it to home base; 3. whether the friendly has enough weight carrying capacity to successfully load the cargo; 4. whether the friendly has enough mission time to load and unload the cargo at home base; and 5. whether the cargo is already allocated to another craft.

The fixed enemy affordance perceptual mechanism indicates the degree to which a fixed ground enemy at a specified location affords being attacked by a friendly craft at a specified location with specified resource levels. The factors that are included in this calculation include: 1. whether the friendly has enough missiles to attack the enemy; 2. whether the friendly is carrying cargo (that should not be jeopardized by such an attack); 3. whether the friendly has enough fuel/time to travel to the enemy and return home; and 4. whether the enemy craft has already been allocated to another friendly craft.

The home base affordance perceptual mechanism indicates the degree to which returning to home base is a desirable action for a friendly craft. This calculation is based on: 1. the number of cargo that need to be unloaded; 2. the number of missiles carried by the friendly; and 3; the amount of fuel the friendly craft has remaining.

The collision affordance perceptual mechanism is charged with detecting potential collisions between friendly craft, and therefore indicating the degree to which evasive action is required.

The mobile enemy affordance mechanism indicates the degree to which an enemy tank or helicopter affords being attacked by a friendly craft. In the case of the friendly craft, this mechanism is quite simple since attacking is maximally afforded when a friendly craft is locked-on by enemy radar, since escaping these enemies is quite difficult to accomplish. In the case of the scout which can sight these enemies at an extended range (unlike the friendlies), the affordance calculation is slightly more complicated.

The friendly-up affordance perceptual processor indicates when a command should be given to a friendly craft to rise above tree level, as craft speed is increased above tree level. This perceptual mechanism operates by detecting when friendly craft are traveling very slowly through heavily forested regions.

Some of these mechanisms would appear to require the integration of information from memory with information from perceptual sources. Therefore, there would seem to be a problem with describing these mechanisms as entirely perceptual, rather than perceptual-cognitive, in nature. One weakness of this work is that the distinction between operations on perceptual and memorial information is not clearly maintained. Rather, it has been assumed that the same type of information integration operations can be used for describing the use of information from both perception and memory. As a result of this assumption, this research may not adequately treat possible information processing limitations related to constrained memory resources. Nevertheless, the assumption that identical operations are used to integrate information from both perceptual and memorial sources will be maintained throughout this work.

The outputs of all these perceptual mechanisms are made available to the object selection mechanism. This mechanism serves to resolve potential conflicts between multiple friendly craft, and to plan future craft activities. The output of this mechanism is a series of commands to the friendly craft action mechanism which simulates the physical interface actions necessary to implement these desired friendly craft activities. Given this understanding of the global organization of the model, the way in which each of these perceptual, selection, and action mechanisms operate can now be described with greater precision.

Description of Mechanisms

The model employs perceptual, selection, and action mechanisms. In abstract terms, perceptual mechanisms provide a mapping from the raw, displayed information into affordance information that is presumably highly relevant for the task of action selection. One metric of the efficiency of these mechanisms is, then, the degree to which they allow for simple action selection mechanisms. A detailed characterization of the perceptual, selection, and action mechanisms is required to demonstrate exactly how the simplicity of selection mechanisms can be achieved by the design of action-oriented perceptual mechanisms.

As suggested by Figure 9, an approximate decomposition of the model structure can be made along subtask lines: the scout waypoint generation subtask, the friendly waypoint generation subtask, and the object (cargo, enemy craft, and home base) selection subtask. The following description of the processing mechanisms follows this decomposition and begins with the scout path planning subtask.

Scout Waypoint Generation Mechanisms

The task of path planning for the scout craft involved selecting a route through the world that maximized the number of cargo and enemy craft discovered with the scout radar. The following general heuristics describe the consensual expert behavior (obtained from post-hoc interviews of skilled crews). The route plan should: 1. Cover as much of the forested regions as possible with scout radar (object densities were higher in forested regions); 2. Avoid traveling through the forests as much as possible (due to slower average speeds through the forest caused by the tree avoidance task); 3. Return to home base about two thirds through the mission to refuel; and 4. Avoid "backtracking" as much as possible (to emphasize previously unsearched regions).

Figure 10 indicates the search paths generated in the same world by the two crews who used manual control of the scout (Crews E and 2). The path generated by Crew 2 that appears in Figure 10 was chosen for illustration because it included a minimum number of deviations to process cargo and enemy craft for this crew. Such deviations obscure the degree to which the route of the scout is indicative of the Crew's desired search path for the scout. The path included for Crew 1 was chosen because it was based on the same world configuration used for Crew 2.

Note the similarity between these paths in that the same general routes were followed by both crews, although the direction of travel along these routes was somewhat different. Both crews visited the same general world locations and tended to stay close to the forest boundaries during travel between these locations. This "boundary hugging" behavior probably resulted from the interaction of the set of competing constraints on path generation discussed above in the characterization of the consensual search strategy. In addition, both crews avoided backtracking, and they both managed to refuel about two thirds through the mission. There is evidence to suggest that the task of planning these routes was completed in no more than two or three seconds, as crews often verbalized the results of their path planning processing within a few seconds after initially viewing the map display at the start of a session.

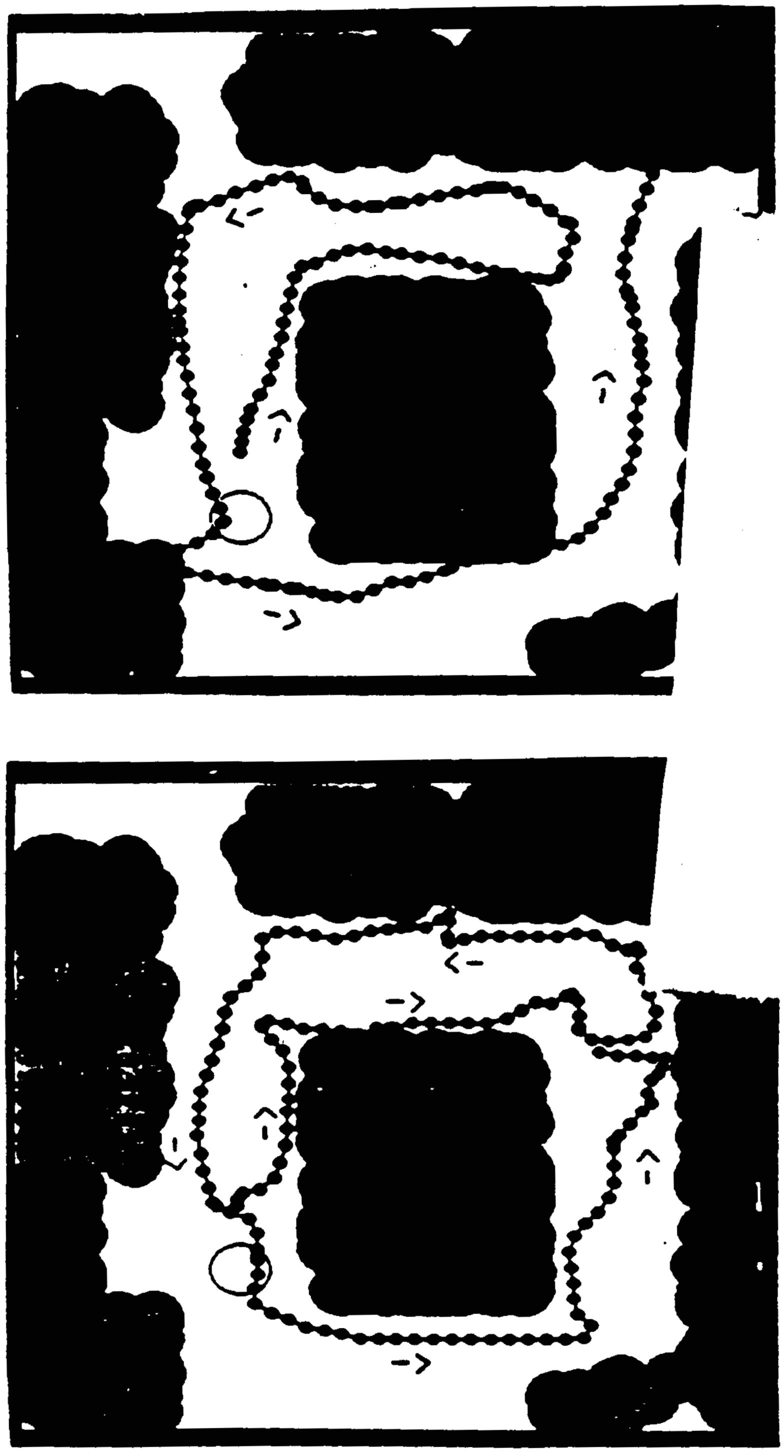

Search paths produced in World 3
by Crew E (top) and Crew 2 (bottom)

Figure 10

As was mentioned in the previous section, the model employs two perceptual affordance mechanisms to transform the displayed information into action oriented information. The scout search affordance mechanism provides a map of the 100 square mile displayed world that indicates the degree to which each world location afforded sighting cargo and enemy craft with radar. The scout locomotion affordance mechanism provides a map of the world that indicates the degree to which each world location afforded speedy locomotion. The combination of the two maps was then used by the search path selection mechanism to provide path planning.

Subjects also received briefing report information at the beginning of each session concerning the probable locations of cargo and enemy craft. It was originally hypothesized that subjects would be sensitive to this information in their selection of search paths. Subsequent analysis, though, did not yield any evidence that subjects used the briefing information, so this information was not included in the model.

The diagrams on the following four pages indicate the search, locomotion, and total affordance maps when applied to each of the four world configurations used in the experiment. The picture in the lower left corner of each figure shows the world as it was presented to subjects on the computer display. The white (non-forested) regions were shown in brown, the moderately dark (lightly forested)

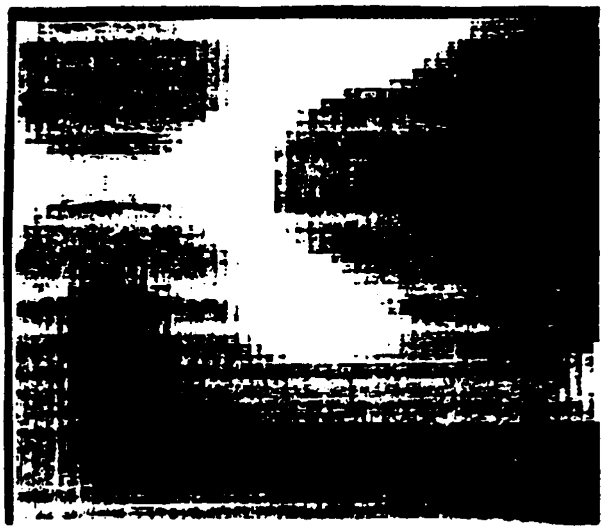

Search Affordance Map

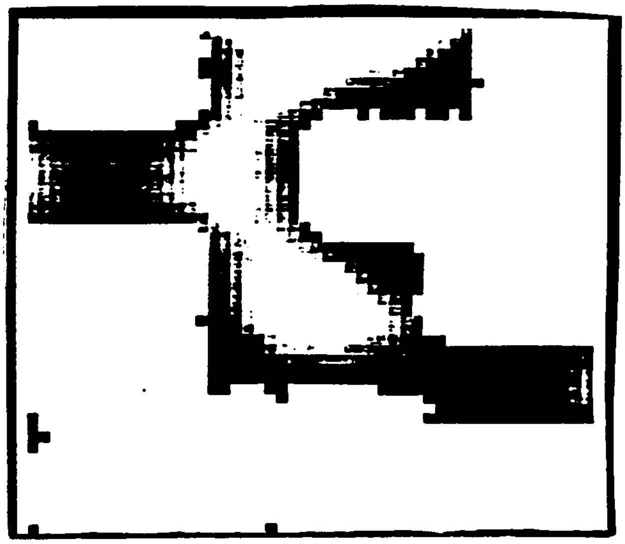

Total Affordance Map

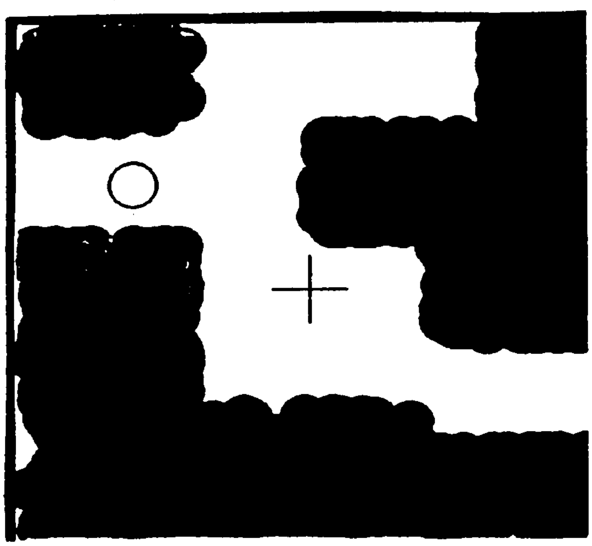

World 1 Configuration

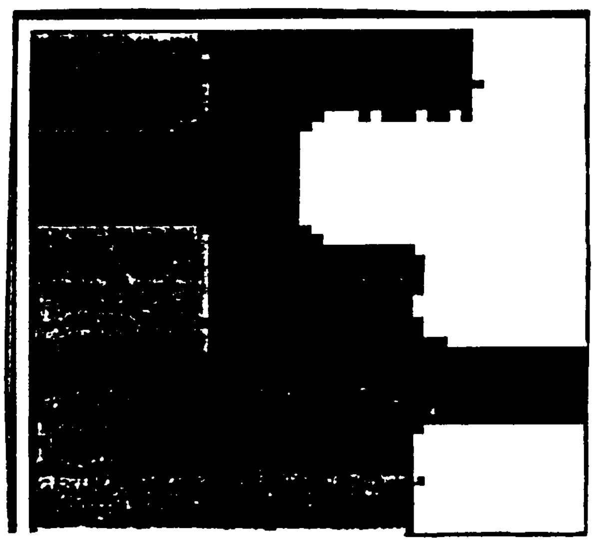

Locomotion Affordance Map

Affordance maps for world 1

Figure 11

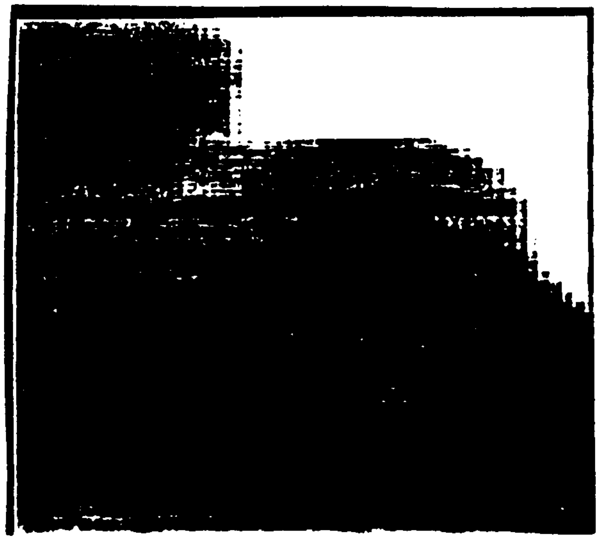

Search Affordance Map

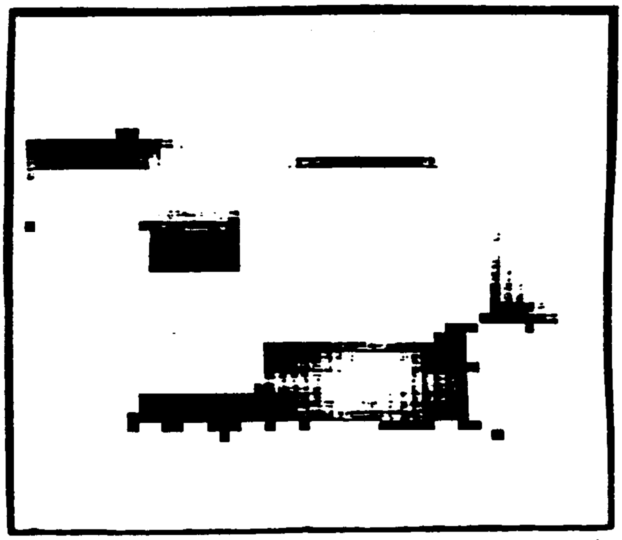

Total Affordance Map

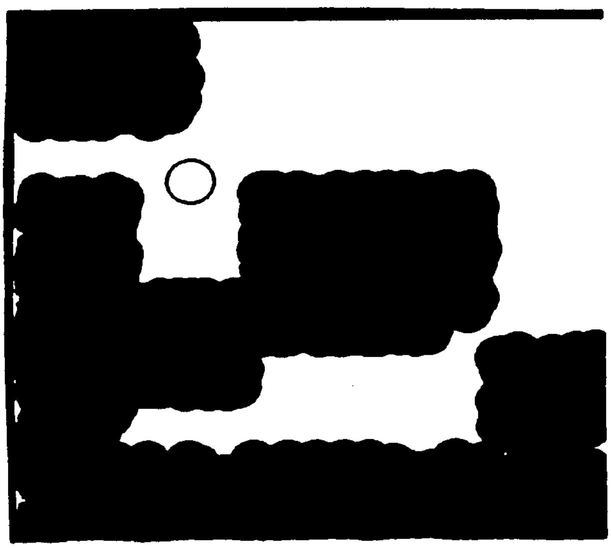

World 2 Configuration

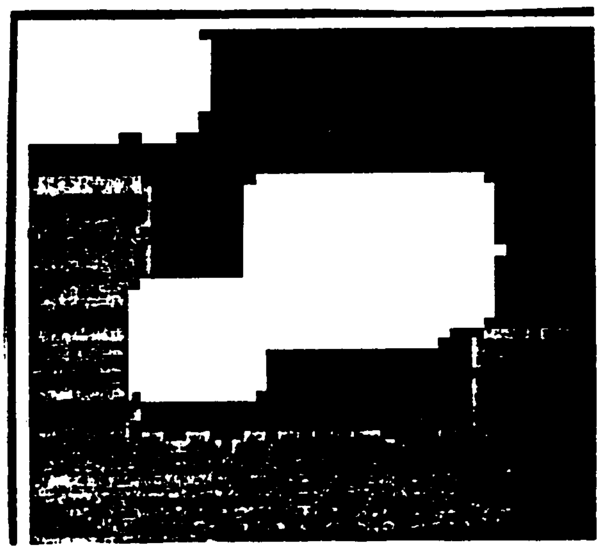

Locomotion Affordance Map

Affordance maps for world 2

Figure 12

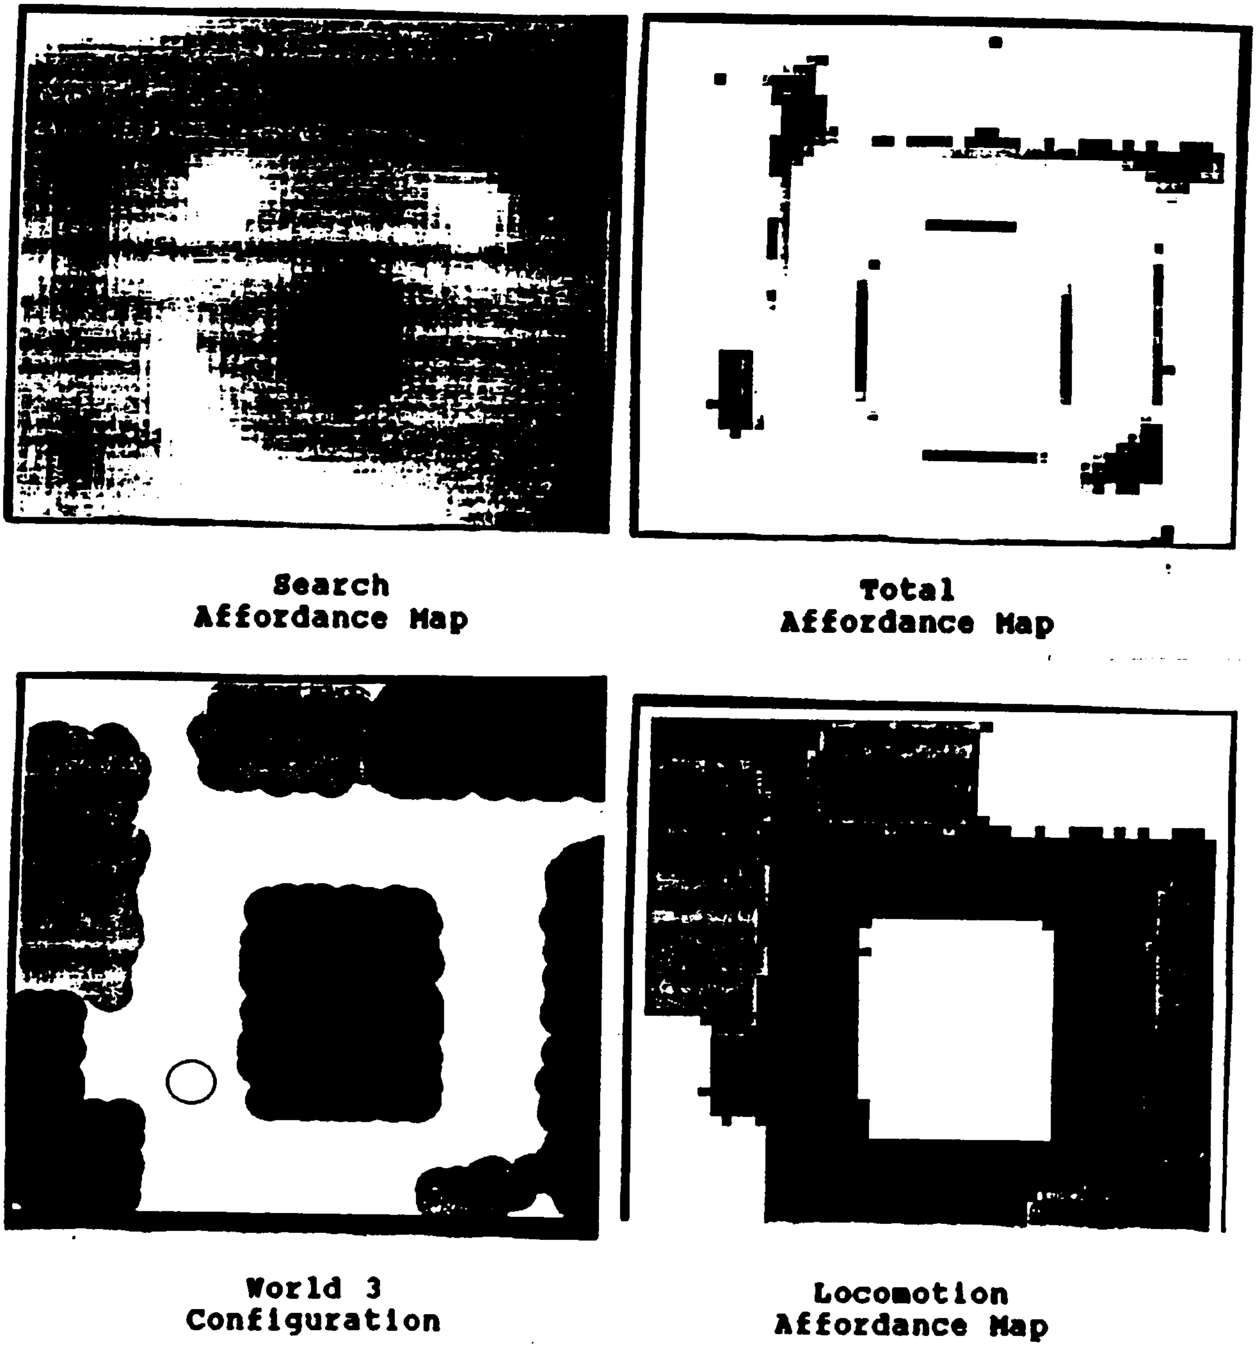

Affordance maps for world 3

Figure 13

ORIGINAL PAGE
BLACK AND WHITE PHOTOGRAPH

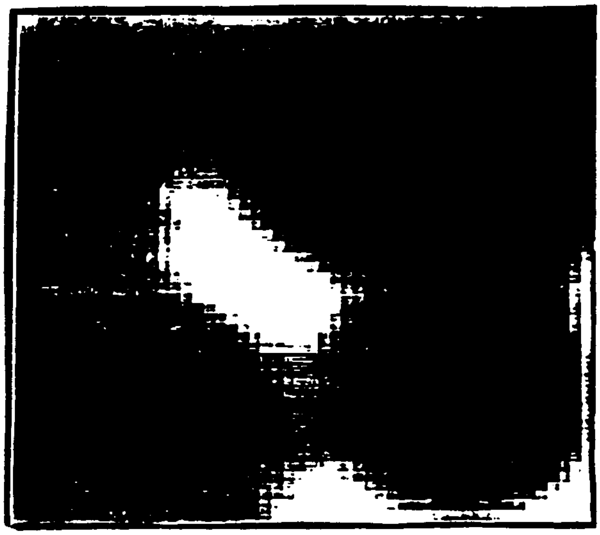

**Search
Affordance Map**

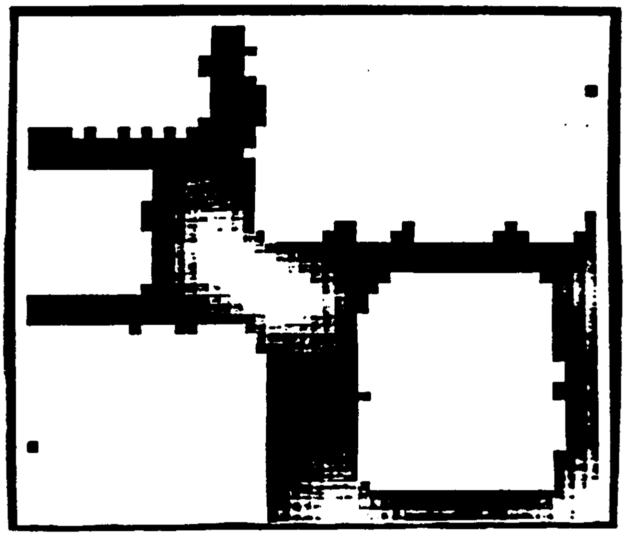

**Total
Affordance Map**

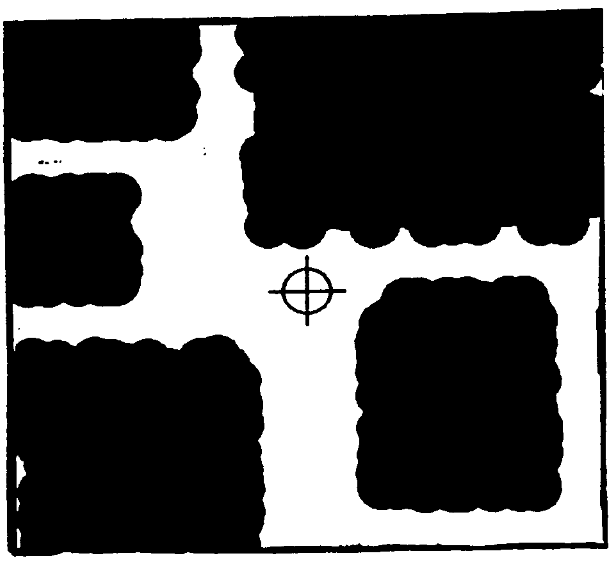

**World 4
Configuration**

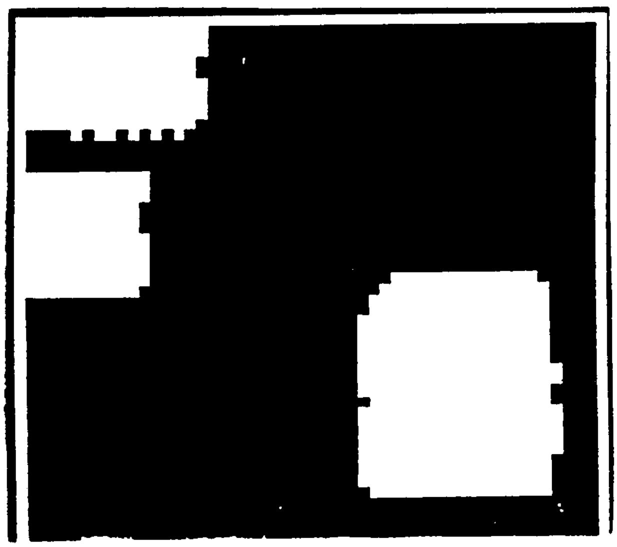

**Locomotion
Affordance Map**

Affordance maps for world 4

Figure 14

regions were shown in light green, and the dark (heavily forested) regions were shown in dark green. The picture in the lower right corner of each figure depicts the locomotion affordance map based on this world configuration that was produced by the scout locomotion affordance mechanism. Darker regions on this map indicate areas of higher locomotion affordance. This is essentially an inverse mapping of the original display since open regions resulted in high search affordances while heavily forested regions resulted in low search affordances.

The picture in the upper left corner of each figure shows the affordance map that was produced by the scout search affordance mechanism. To construct this map, a four dimensional vector was associated with each world location to indicate the percentage of area that would be covered by scout radar that was open region, lightly forested region, heavily forested region, and beyond the boundary of the simulated world. For each point, the inner product of this vector and a search affordance vector is taken to determine the total search affordance of the world location. The search affordance vector was the same for each world location and indicated the perceived density of objects within each of the four regions. Darker points on this map indicate a higher coverage of forested regions with scout radar centered at the point.

The result of combining (and rescaling for clarity) these two affordance maps is shown in the picture in the upper right corner of each figure. This map indicates the degree to which each world location affords both actions of locomoting and sighting objects collectively. This map was produced by adding the affordance values of the two previous maps under a constant weighting scheme. This process of map superposition as a method of value integration is similar to one discussed by McHarg (1971) for ecological planning of highway routes.

Figure 13 shows the maps produced when applied to the same world shown in Figure 10. Notice that the peak (darkest) areas in the total affordance map in Figure 13 roughly correspond to the regions covered by the two crews' paths. This total affordance map is the information upon which the scout waypoint selection mechanism operates. The scout waypoint selection mechanism operates by considering only the peak areas in the total affordance map as candidate search waypoints for the scout. At the beginning of the mission, a search and evaluation mechanism cycles through these peaks to select a search plan for the mission that satisfies two path planning criteria in addition to the search and locomotion criteria provided by the perceptual mechanisms. These two additional planning constraints ensure that the selected path minimizes backtracking through previously searched regions and passes through home base in the middle third of the mission for refueling. The path plan that visits all the peaks and maximizes searched area subject to the backtracking and refueling constraints is chosen. Although the need for generation and evaluation of action alternatives has not been eliminated, the perceptual affordance mechanisms radically decrease the size of the space of action alternatives that must be considered by the evaluation process by providing a small set of peaks (typically about ten) as waypoints. Routes among these waypoints are much more quickly evaluated than would be routes among a set of candidate waypoints provided by a discretization of the world without regard to search and locomotion affordance information.

Friendly Craft Waypoint Generation Mechanisms

The mechanisms for friendly craft waypoint generation operate in a similar fashion by constructing independent affordance maps and then combining them to produce a total affordance map. Friendly craft are given search commands by the specification of a waypoint on the map display via the manipulation of the map cursor and typing a search command on the text editor. The four perceptual affordance mechanisms that provide the information on which this task is performed include the friendly search affordance, friendly collision affordance, friendly-scout cooperation affordance, and friendly range affordance processors. Recall that the scout search affordance processor determined the search affordance of a world location by summing the search affordances of all the world points that were within scout radar radius of the location. The friendly search affordance processor also operates by summing search affordances, although the dimension along which the summation is performed is different.

Since assigning a friendly craft waypoint determines a linear path of travel from the initial craft position to the waypoint, it is clear that waypoint selection is actually the selection of a linear route of travel, where the search width of this route is the radar radius of the friendly craft (0.4 miles). The individual affordance values that are summed to determine the search affordance of a given world location are the affordance values of the world points contained within the 0.4 mile wide rectangular region traced out by the craft's path to the waypoint. Therefore, unlike the map produced for the scout, the affordance map used for friendly craft waypoint selection is a function of the initial craft position, and therefore must be created each time the waypoint selection task is initiated.

A second affordance map is produced by the friendly collision affordance perceptual mechanism. This mechanism produces an affordance map that is binary valued. Areas of potential of collision with other craft (determined by the planned trajectories of the other craft) receive zero affordance value, and all other areas receive unit value.

In the map combination process, a multiplicative rule is used to integrate this map with the others, so regions of possible collision are excluded from consideration.

A third affordance map for the task of friendly craft waypoint selection is produced by the friendly-scout cooperation affordance mechanism. To describe crews who appear to spatially coordinate the scout and friendly craft, this mechanism can be used to provide an identification of those world locations that, given the planned scout path, afford simultaneous searching by the scout and friendly craft. The specified area is the locus of world points that can be simultaneously covered by both friendly craft and scout radar. Like the collision affordance map, this map is also binary valued, although the affordance values on this map are summed with the other maps, rather than multiplied, so that points where no coordination is achieved receive no additional value but are not excluded from consideration.

The final perceptual mechanism used for friendly craft waypoint selection is the friendly range affordance mechanism. This mechanism also produces a binary valued affordance map, where unit value is assigned to those points that are within range of the friendly craft, and zero value is assigned to those points that are outside of the craft's range. In the condition where a craft has enough fuel to last until the end of the thirty minute mission, all world points are considered to be in range. When this condition is not met, a world location is determined to be in range if the craft can reach the world location and also travel from the location to home base to refuel. The locus of all such points defines an ellipse with foci at the friendly craft and at home base. The eccentricity of the ellipse is determined by the relationship between the craft's fuel level and its distance from home base. The affordance values from this map are multiplied with the values from the other three maps so that world locations that are outside of the craft's range are excluded from consideration.

The friendly waypoint selection processor operates on the combination of these maps under the additive or multiplicative rules described above. As was the case with the scout search path selection mechanism, the waypoint selection mechanism is quite simple. The mechanism operates by selecting as a waypoint the world location from the total affordance map with maximum affordance value. One result of the way that the affordance maps are constructed is that, given that world border locations are within range and not potential collision points, friendly waypoints selected under the method described above will always be world border locations. This result arises due to the fact that the search affordance function is always non-decreasing as the distance from the friendly craft increases, due to the summation of affordances along the linear routes.

As will be described in the following chapter, the waypoints generated by crews do not follow this pattern, as subjects were reluctant, for some as yet unidentified reason, to send friendly craft close to the borders of the simulated world. Therefore, an ad-hoc constraint placed on the selection of friendly craft waypoints was that world locations that were close to the world borders were excluded from consideration. If the reason for this behavior could be identified, or even soundly hypothesized, a perceptual affordance mechanism could be easily constructed that would produce a binary valued map that could be multiplied with the affordance values of the other four maps to exclude these extreme locations.

World Object Planning Mechanisms

The two search oriented components of the modeling structure have now been described. The remaining component is the set of perceptual and selection mechanisms that describe the method by which goal objects (cargo, enemy craft, and home base) are assigned to friendly craft, and the action mechanisms that allow these actions to be implemented by simulating the crews' physical interaction with the interface.

The perceptual affordance mechanisms devoted to the object selection task include the cargo affordance, fixed enemy affordance, home affordance, collision affordance, mobile enemy affordance, and friendly-up affordance perceptual processors. These mechanisms determine the degree to which particular objects afford actions of various types to friendly craft depending upon the particular friendly craft resources. With the exception of home base, each type of object defines one type of affordance, since only one type of action is associated with each object type (i.e. cargo afford loading, enemies afford attacking, potential collisions afford evasive action, and friendlies moving slowly to avoid trees afford being sent above tree level). Home base, on the other hand, affords three types of actions: refueling, resupplying missiles, and unloading cargo. A high affordance for any one of these actions can be sufficient reason to be sent to home base.

As was the case with the perceptual mechanisms described earlier, one role of the object perceptual affordance mechanisms is to simplify the problem of the object action selection mechanism. The task of determining an efficient allocation of craft to world objects can be viewed as a complex optimization problem. One way that a sensitivity to affordances can reduce the complexity of an optimization problem was discussed above in reference to scout waypoint selection, where use of affordance information resulted in a drastic reduction in the size of the problem space to be searched. A sensitivity to affordances can also reduce the complexity of an optimization problem by decomposing that problem into a large number of relatively simple, independent sub-problems. Each of the sub-problems is the assessment of an action-environment relation, and will reflect both the desirability of a particular action as well as the degree to which it can effectively be performed given the specifics of the environmental situation.

Problems of optimization almost invariably involve a minimization of resources to achieve a fixed objective, or a maximization of an objective function given fixed resources. That is, they require an integration of costs and benefits at some point in the computation using a common scale of measurement. Affordances serve a similar function. Since they are measures of the desirability of taking a specified action in a specified environment, they provide implicit tradeoff evaluations that result in a one-dimensional measure of both the benefits and the costs of the action. They do this by reflecting both the desirability of the action (the benefits) and the degree to which the environment allows for effective performance of the action (the costs). To suggest that the perceptual mechanisms are attuned to affordances, then, is to imply that the necessity for explicit trade-off comparisons is reduced, due to the fact that the perceptual mechanisms perform an implicit cost-benefit evaluation. This is one explanation of how attunement to affordances can reduce the burdens on an information processing system that is charged with a task that requires the comparison of costs and benefits in the selection of action.

The object affordance perceptual mechanisms are concrete examples of this principle in operation. The cargo affordance mechanism provides a scalar value that is indicative of both the desirability of loading a particular cargo and the ease with which this task can be performed. The affordance is actually a scalar measure that reflects a combination of both these dimensions. These dimensions are reflected in the cargo affordance mechanism by the fact that the cargo affordance is increased with decreased distance between the friendly craft and the cargo. One of the costs involved in loading a cargo is the time spent by the friendly to perform this task, and this time decreases with decreased distance to the cargo. Another potential cost is losing the friendly craft by running out of fuel on the way to load or unload a cargo, so the cargo affordance mechanism provides a zero value in this condition. On the benefit side, the payoff for the cargo can only be achieved if the craft has enough time to return the cargo home, and enough missiles to defend the cargo (and friendly craft) from being destroyed, therefore the affordance mechanism is also sensitive to these considerations. Of course, the decision as to which dimensions of an action are to be considered costs and which factors are benefits does not seem to be an objective property of the action; rather it appears to depend upon the goal structure that the action serves and other contextual factors.

The fixed enemy affordance mechanism operates in a similar fashion. Considerations to which this mechanism is sensitive are the distance between the friendly craft and the fixed enemy, the number of missiles possessed by the friendly craft, the number of cargo carried by the friendly craft that might be sacrificed due to an unsuccessful attack action, the friendly craft's fuel level that determines whether the enemy is in range, and the mission time remaining in the case that the friendly craft will have to return home before the end of the mission. The affordance mechanism produces a scalar measure of the combination of this entire set of factors. The detailed operation of this and each of the mechanisms discussed below will be described in greater detail in the following chapter.

The mobile enemy affordance mechanism is much more simple due to the fact that, at the point when attacking mobile enemies is a feasible action for a friendly craft, the mobile enemy is endangering the existence of the friendly craft. Thus, unless the attack action is taken, the friendly craft will typically be destroyed. Therefore, a maximum affordance value is associated with lock-ons with mobile enemy craft, otherwise a zero value is assigned to this affordance.

The home base affordance mechanism is more complicated due to the fact that home base affords three types of action. The craft may return home to refuel, to resupply with missiles, to unload cargo, or any combination of these three activities. Therefore, separate affordance calculations are made for each of the three possible actions, and the resultant affordance values are then combined into a single measure by the use of a rule that contains both additive and maximization components. The additive component is due to the fact that the three affordances can combine in a way that, although none of the affordance values on its own is high enough to suggest that the craft should be sent home, going home still might be desired due to a combination of needs. The maximization operation arises because any one of the three affordance dimensions, though, is sufficient for requiring a home action, regardless of the values along the other two dimensions.

It would appear at least theoretically possible to eliminate the maximization component of this mechanism with an appropriately designed additive operation. An additive rule can be used to allow for a high input value on one information dimension to be sufficient, on its own, for the generation of a high output value. The maximization operation was used, though, because it provides an implicit thresholding function that cannot be performed by an additive operation alone. The reason that thresholding was desired was to allow a very high value on one dimension to null the values on the other dimensions in the output calculation. Consider, for example, the situation where it was absolutely necessary that a craft return home to refuel, resulting in a maximum home affordance value. The use of a maximization operation in this situation would yield the same (maximum) affordance value regardless of the values on the other input dimensions. Since the lack of fuel made it absolutely necessary that the craft return home, an additional need to unload a piece of cargo, for example, could make it no "more" necessary that the craft return home.

The final two perceptual affordance mechanisms help determine when two additional friendly actions should occur, although these actions do not have an associated world object as do the actions discussed above. One of these is the collision avoidance affordance mechanism, which is sensitive to potential collisions between friendly craft. A high affordance value indicates that a potential collision is imminent, and suggests that a collision avoidance action should be implemented. This mechanism operates by performing an extrapolation of each craft's position to identify potential collisions. The final affordance mechanism is sensitive to situations where the friendly craft is traveling slowly due to the fact that it is below tree level and within a forested region. In this situation, the friendly craft should be given a command to raise its altitude to above tree level. This situation is not highly salient perceptually, as it requires repeated observations of the craft's motion on the map display (to generate a speed estimate) to be identified. Similarly, the affordance mechanism that is sensitive to this situation is progressive: the affordance value is slowly increased by small amounts during the time interval in which the situation exists. This feature is implemented by adding a small number to this affordance value each time the affordance values are updated when the model "looks" at the world display.

These five perceptual affordance mechanisms supply the information that is used by the object selection mechanism. This mechanism operates upon the entire set of affordances between all craft-action pairs that are made available by the current, and predicted, world state and determines the desired current and planned craft actions. For the scout and each of the four friendly craft, the affordance values for each action the world currently makes available are calculated based on current world state, the world state that is predicted to exist at the end of the craft's current activity, and the world state that is predicted to exist at the end of the craft's currently planned activity. A given cargo that is discovered, for example, can be allocated to any one of the friendly craft at any one of three time points (currently, immediately following the current action, or immediately following the planned action).

The allocation mechanism operates in the following way. Let i, (i=1,5) range over the five craft, where i=5 designates the scout craft. The planning horizon for determining craft activity is assumed to be three actions. (Results of a sensitivity analysis that suggested that the assumption of a three action planning horizon was required to mimic crew behavior are discussed in the next chapter). Let j, (j=1,3) range over the three time slots in the planning horizon for each craft: j=1 is the current action; j=2 is the first planned action; and j=3 is the second planned action. Let k, (k=1,N) designate the actions that are made available to a craft at a given time, where N is determined by the current and predicted world state. Each k indicates a particular action available to a friendly craft, where k=1 might indicate loading cargo #3, k=2 might indicate loading cargo #12, k=3 might indicate attacking fixed enemy #7, and k=4 might indicate going home.

Now, let A(i,j,k) be the affordance (provided by the perceptual processors) for assigning friendly(i) at time slot(j) to object(k). When j=1, the A(i,j,k) are calculated directly from the current world state and the current resource levels for each of the craft. If a craft is currently performing an action for which a termination time can be estimated (all actions except search actions which continue indefinitely), the A(i,2,k) are calculated based on the world state and the craft resource levels that are predicted to exist at this time. The predictions are made with a set of simple difference equations that describe system dynamics. The equations describe how craft the position and resource levels change over time. In the case where a given friendly has been scheduled to perform an action after the current action, and the scheduled action also has a predictable termination time, the A(i,3,k) are calculated based on the world state and craft resource levels that are predicted to exist at the initiation of the third time slot in the planning horizon.

Each craft is treated individually when making these predictions. That is, the mechanism serially cycles through each of the craft independently when advancing craft to future time points. Thus, a comprehensive prediction involving all five craft is not used for planning future craft actions. Thus, the mechanism is not able to plan future, coordinated activity among the multiple craft. Rather, the mechanism can only plan future actions that do not involve coordination. The ability to generate co-ordinated actions is restricted to actions taken in the current world state. That is, the perceptual mechanisms that are used to detect coordination affordances are not used for planning future actions, since the estimates of future world states are made for the craft individually.

Since the possibility exists that optimal performance requires planning coordinated actions, this mechanism should be considered as a heuristic, rather than an algorithm.

As mentioned previously, the function performed by this planning mechanism is assumed to be performed by visual imagery. The reason that the mechanism was not designed to be able to generate planned, coordinated actions concerns an assumption about how imagery might be limited in its ability to provide veridical simulations of future world states. Although no empirical evidence has been cited to support this claim, it is assumed that the human's ability to advance the world state via visual imagery is restricted to individual craft. When imaging the future trajectory of an individual craft, it is assumed that crews cannot image the trajectories of the other craft that appear peripheral to the craft which is the focus of attention. Since the perceptual mechanisms sensitive to coordination affordances require information pertaining to the simultaneous location of multiple craft, the operation of these mechanisms is assumed to be restricted to processing the currently displayed world state only. While it may be the case that humans are capable of simultaneously imaging the future states of multiple objects (e.g., highly skilled chess players imaging the future location of multiple chess pieces), this capability is not assumed to be operative in this experimental task.

After these predictions have been made, the three dimensional array A(i,j,k) indicates the affordance values for each of the five craft at each of the three time slots for each of the actions that the world makes available. A comparison of these values is required due to the fact that objects supporting world actions (cargo and enemy craft) are "consumable" and cannot be allocated to more than one friendly time slot. The action selection mechanism operates by assigning the most appropriate action to each friendly time slot under the constraint that multiple actions requiring the same world object are not allowed. That is, cargo #3, for example, cannot be allocated to both friendly #1 and friendly #2, nor can it be allocated to both friendly #1's current time slot and friendly #1's second time slot.

For the object action selection mechanism, alternatives need to be compared so that the same object is not allocated to multiple craft. These comparisons are necessary due to the fact that the problem decomposition used to create the perceptual affordance mechanisms did not result in completely independent sub-problems. The selection mechanism is required, then, to deal with the interactions between the solutions to the sub-problems caused by the fact that the problem decomposition was not entirely clean.

To satisfy the exclusivity constraint, the following heuristic is used to determine the allocation of consumable objects based on the affordance values in array A(i,j,k).

Simply said, the heuristic tries to assign the "best" friendly (as determined by its location and resource levels) to each consumable object. For example, let k=1 be a consumable object that should be assigned to a craft. Given the current world state and the current and planned activities for the five craft, affordance values for this cargo, A(i,j,1), are determined for each craft(i) and time slot(j).

The heuristic makes use of the following policy to determine to which craft time slot each consumable object should be assigned. The heuristic is defined below.

1. Let OMAX(ij) = m, if Max(A(i,j,k), {k=1,n}) = A(i,j,m)

   That is, action(m) is the best action for friendly(i) in time slot(j).

2. Let FMAX(k) be the ordered pair (f,t) if:

   Max(A(i,j,k), {i=1,4}, {j=1,3}) = A(f,t,k)

   That is, the best craft time slot for action(k) is for friendly(f) at time slot(t).

3. Allocate action(k) to friendly(i) at time slot(j) if

   1. OMAX(i,j) = k and;

   2. FMAX(k) = (i,j)

   That is, allocate action(k) to friendly craft(i) at time slot(j) if action(k) is the most highly afforded action for friendly(i) at time slot(j), when considered across all available actions; <u>and</u> friendly(i) at time slot(j) is the highest affordance for action(j) considering all friendlies and time slots.

Allocation of consumable objects to craft is quite simple under this heuristic. The heuristic is motivated by considering how a human supervisor might allocated tasks to a staff of people (the five craft) who vary in their abilities to perform those tasks (craft locations and resource levels). The goal is to "find the best person for the job", and to simultaneously "find the best job for each person". Now, when the allocations made under these two principles agree, the allocation task is easy. For example, consider when a new task is added to the job queue and the hypothetical human supervisor must assign it one member of the support staff. The people comprising the staff may differ in their ability to perform the task so the supervisor would like to find the person who can perform the task most efficiently. On the other hand, the supervisor does not want to assign the task to a person who is better used to perform a different task. Sally, for example, might be the most highly skilled computer programmer on the staff. Assuming that she is not already performing other programming tasks of more importance than the new task, the new task would most likely be allocated to her.

Problems can arise, of course, when the most highly skilled person for the task is unavailable. In this situation, the heuristic described above leaves this task unallocated. To deal with this situation, the human supervisor might try to find the next most competent staff member. This operation is accomplished by adding the following feature to the allocation heuristic: when a craft has a more highly afforded action in a given time slot than the new object that needs to be allocated, the affordance for the object is nulled because the craft has a more highly afforded action. Returning to the example of the human supervisor, this rule says that when a task is best performed by Sally, but Sally is busy with more important tasks, the affordance that Sally feels for the new task, even though she is the most competent performer of the task, is essentially zero. In this way, Sally no longer competes for the new task, and another staff member can assume the role of the most competent performer due to Sally's unavailability. The friendly craft are allocated to consumable world objects (tasks) using the same scheme. It should be remembered, though, that the competition for objects is not only between craft but also between the time slots for each craft.

For example, consider a situation where craft 1 is on its way to home base to refuel and resupply with missiles. Assume a cargo is then discovered and the affordances for this cargo for each of the craft are identified. Craft 1 might have the highest cargo affordance value among the five craft, yet it might have an even higher home affordance value due to its need for missiles. In this situation, the heuristic described above would null craft 1's current affordance value for the cargo so that other craft and other craft 1 time slots can compete for the cargo. For example, craft 2 might now have the highest cargo affordance value so the cargo would be allocated to craft 2. On the other hand, it might be the case that the highest cargo affordance value would be possessed by craft 1's second time slot, indicating that craft 1 should load the cargo after it is refueled and resupplied with missiles at home base.

In a different scenario, craft 1 might be locked-on by an enemy helicopter during its trip home to refuel. Even though it has a high home affordance value, the affordance value for attacking the enemy helicopter would be higher. In this case, the home action for craft 1 would be interrupted and the home action would be rescheduled to the second time slot (after attacking the helicopter).

The frequency with which the planning heuristic described above results in nulling affordance values and interrupting ongoing activity is determined by the complexity of the world situation to which the perceptual mechanisms are applied. Early in the mission when no cargo or enemy craft have been sighted and all craft have high resource levels, the affordance matrix A(i,j,k) would simply indicate a high current affordance for searching for each of the craft. As a result, the entries in the matrix corresponding to planned actions would be zeros, since search actions continue indefinitely. In the middle of the mission, on the other hand, many cargo and enemies may have been sighted and some of the friendly craft may need to return home for resources. In this situation, the affordance matrix A(i,j,k) may be nearly full of non-zero values, reflecting the complexity of the task environment. Although it may be possible to hypothesize a relationship between the situation-dependent complexity of the manipulations on the affordance matrix and dynamic measures of mental workload, this issue was not addressed in this research.

To illustrate how the affordance values in the matrix A(i,j,k) depend upon the current world situation, refer to Figure 15 on the following page. This figure indicates a possible world situation and the corresponding approximate affordance values that would be generated by the perceptual mechanisms. The array of distributions at the top of the figure is a graphical representation of the affordance matrix A(i,j,k). Each row (i) represents a craft, with the friendly craft in the top four rows and the scout in the bottom row. Each column (j) represents a time slot in the planning horizon for each craft. From left to right, the columns indicate the first, second, and third actions in the planning horizon. Each vertical bar (k) within each distribution indicates the affordance value for candidate craft action. From left to right within each distribution, the candidate actions are going home (H), searching (S), loading a particular cargo (C1,C2,C3), attacking a red enemy craft helicopter (R), attacking an orange enemy tank (O), and attacking a yellow fixed enemy emplacement (Y). For clarity, the vertical bars associated with any additional

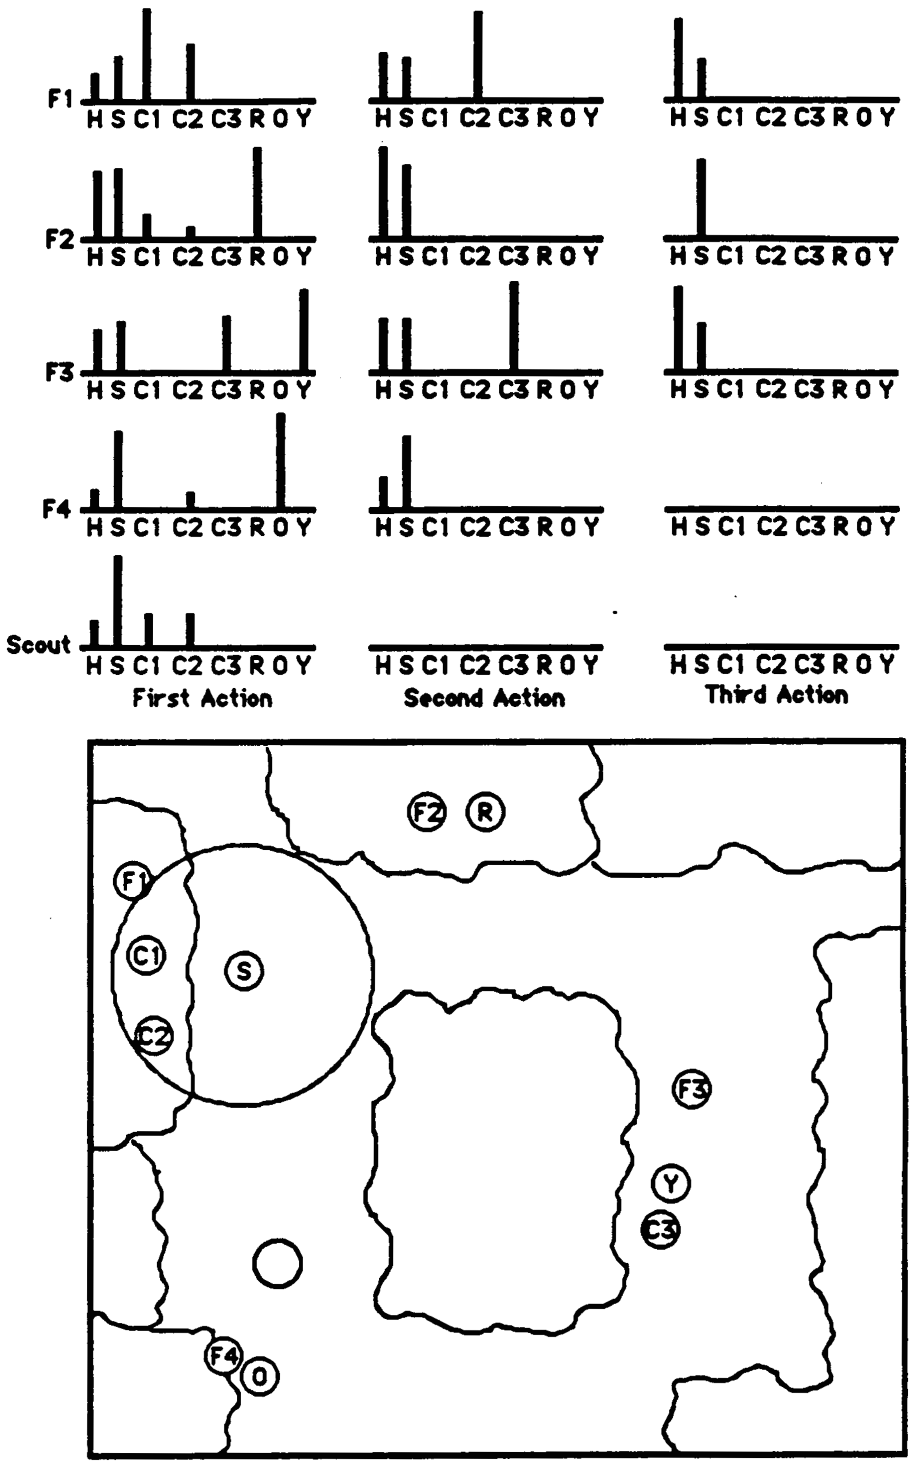

Figure 15. At top, the affordance distributions for each craft indicative of the world situation shown below. F=Friendly craft, S=Scout, C=Cargo, R=Enemy helicopter, O=Enemy tank, Y=Fixed enemy emplacement. Large circle indicates scout radar range.

cargo and enemy craft do not appear in the diagram. Recall that there were actually twelve pieces of cargo, and six enemy targets of each type. A computer generated dynamic display of the type shown in Figure 15 was used during the modeling process to observe the model's internal processing operations.

In the situation shown in the bottom of the figure, the scout is in the northwest region as indicated by the small circle labeled "S". Assume that the scout has just been flown above tree level, thereby sighting cargo C1 and C2 which are within the 1.5 mile radar radius of the scout. For clarity, area covered by scout radar is indicated by the large circle centered at the scout location. A circle indicating scout radar coverage was not included on the world display used by human crews. Another previously sighted cargo, cargo C3, is shown in the southeast region of the world. The four friendly craft are indicated by small circles labeled F1-F4. The circles labeled "R", "O", and "Y" indicate a red enemy helicopter, an orange enemy tank, and a yellow fixed enemy emplacement, respectively. The slightly larger, empty circle in the southwest region of the world indicates home base.

The affordance values represented by the heights of the vertical bars in the array of distributions are specific to a given craft and a given slot in the craft's planning horizon. For example, the affordance values in the distribution at the upper left of the figure indicate the appropriateness of assigning each of the candidate actions to friendly #1's first time slot. As the affordance values indicate, the most appropriate current action for friendly #1 is to load cargo C1. Assume, also, that this affordance value is the highest affordance value for loading cargo C1 when compared to all other craft-time slot pairs. This cargo would then be allocated to friendly #1's first time slot and the affordance values for all other craft-time slot pairs for loading this cargo would be set to zero.

To show, though, how different craft might have different affordance values for the same action, the C1 affordance values for each of the other four craft are shown in the remaining rows in the first column of distributions. These lower cargo affordance values are only meant to be indicative of the greater distances between these craft and cargo C1, since in this simplified example it is assumed that all craft have enough fuel and weight carrying capacity to load any of the cargo in the world. The affordance values shown for loading cargo C1 in the figure, then, are those that would occur just prior to friendly #1 winning the competition for cargo C1. Just after this time, the affordance values for loading cargo C1 for each of the other four craft would be set to zero.

The distribution in the second column for friendly #1 indicates the affordance values for each of the candidate actions that are predicted to exist at the termination of the craft's current action. Since the current action is to load cargo C1, these affordance values indicate the attractiveness of each of the candidate actions given the world state that is expected to exist just after friendly #1 loads cargo C1. The affordance value for the home action in this distribution is slightly higher, reflecting the fact that friendly #1 will then have a cargo (C1) that needs to be unloaded at home base. Of even higher value, though, is the affordance value for loading cargo C2, reflecting the small distance between friendly craft #1 and cargo C2 that is predicted to exist when the craft finishes loading cargo C1. It is assumed that friendly #1 has won the competition for cargo C2 as well, therefore the affordance values for the other craft for loading this cargo have already been set to zero.

The distribution in the third column for friendly #1 indicates the affordance values for each of the candidate actions that are predicted to exist at the termination of the craft's second action. Since friendly #1's second action will be to load cargo C2, the craft will now then have two cargo that need to be unloaded at home base. The high affordance value for the home action in this distribution reflects this fact. Friendly #1's third action would therefore be to go to home base to unload the cargo and resupply with fuel and missiles.

The distributions for each of the other craft should be interpreted in the same fashion. Friendly #2 is locked-on by a red enemy helicopter, and the distributions indicate that this craft should attack the helicopter and then proceed to home base. This situation might result if friendly #2's attack was expected to deplete its missile supply to a dangerously low level. After going home, the third distribution indicates that friendly #2 should be given a search action.

Friendly #3 is scheduled to attack a yellow fixed ground enemy emplacement, load cargo C3, and then travel to home base, presumably to unload the cargo and resupply with missiles and fuel. Friendly #4 is attacking an orange enemy tank nearby home base, and is then scheduled to search. Perhaps friendly #4 has just resupplied with missiles at home base and the encounter with the tank is not expected to decrease its missile supply to a dangerously low level. The scout is currently searching and no future scout actions have been planned.

Once the determination of desired scout and friendly craft actions is made with this planning procedure operating upon information provided by the perceptual affordance mechanisms, control actions need to be taken at the interface to implement these actions. The mechanisms responsible for organizing and producing these actions are described in the following section.

<u>Action Mechanisms</u>

The model's action mechanisms responsible for simulating the crews' physical interaction with the interface are shown in Figure 16 on the following page. The action component of the model is a roughly hierarchical arrangement of mechanisms. Although the mechanisms at the top of the diagram are more abstract in the sense that they are distant from the bottom level mechanisms that actually simulate physical activity, they exist in their own right, and do not simply represent higher-level descriptions of the organization of the activities of the low-level mechanisms. Rather, they are responsible for managing the flow of information to the physical action mechanisms and for coordinating physical activities that require multiple physical action mechanisms.

Hierarchically arranged control mechanisms have been used by many behavioral scientists to describe skilled human action (Jagacinski, Plamondon, and Miller, 1987; Pew, 1984; Mackay, 1984; Harvey and Greer, 1980; and Norman, 1981). One major assumption of this view is that action control is distributed in the form of "schemas" (Norman), "motor programs" (Rosenbaum and Saltzman, 1984), or "activity modules" (Jagacinski et. al.). While each of these constructs is slightly different, they all share the property that they possess a degree of autonomous control of action which does not have to be directly controlled at higher levels. Each

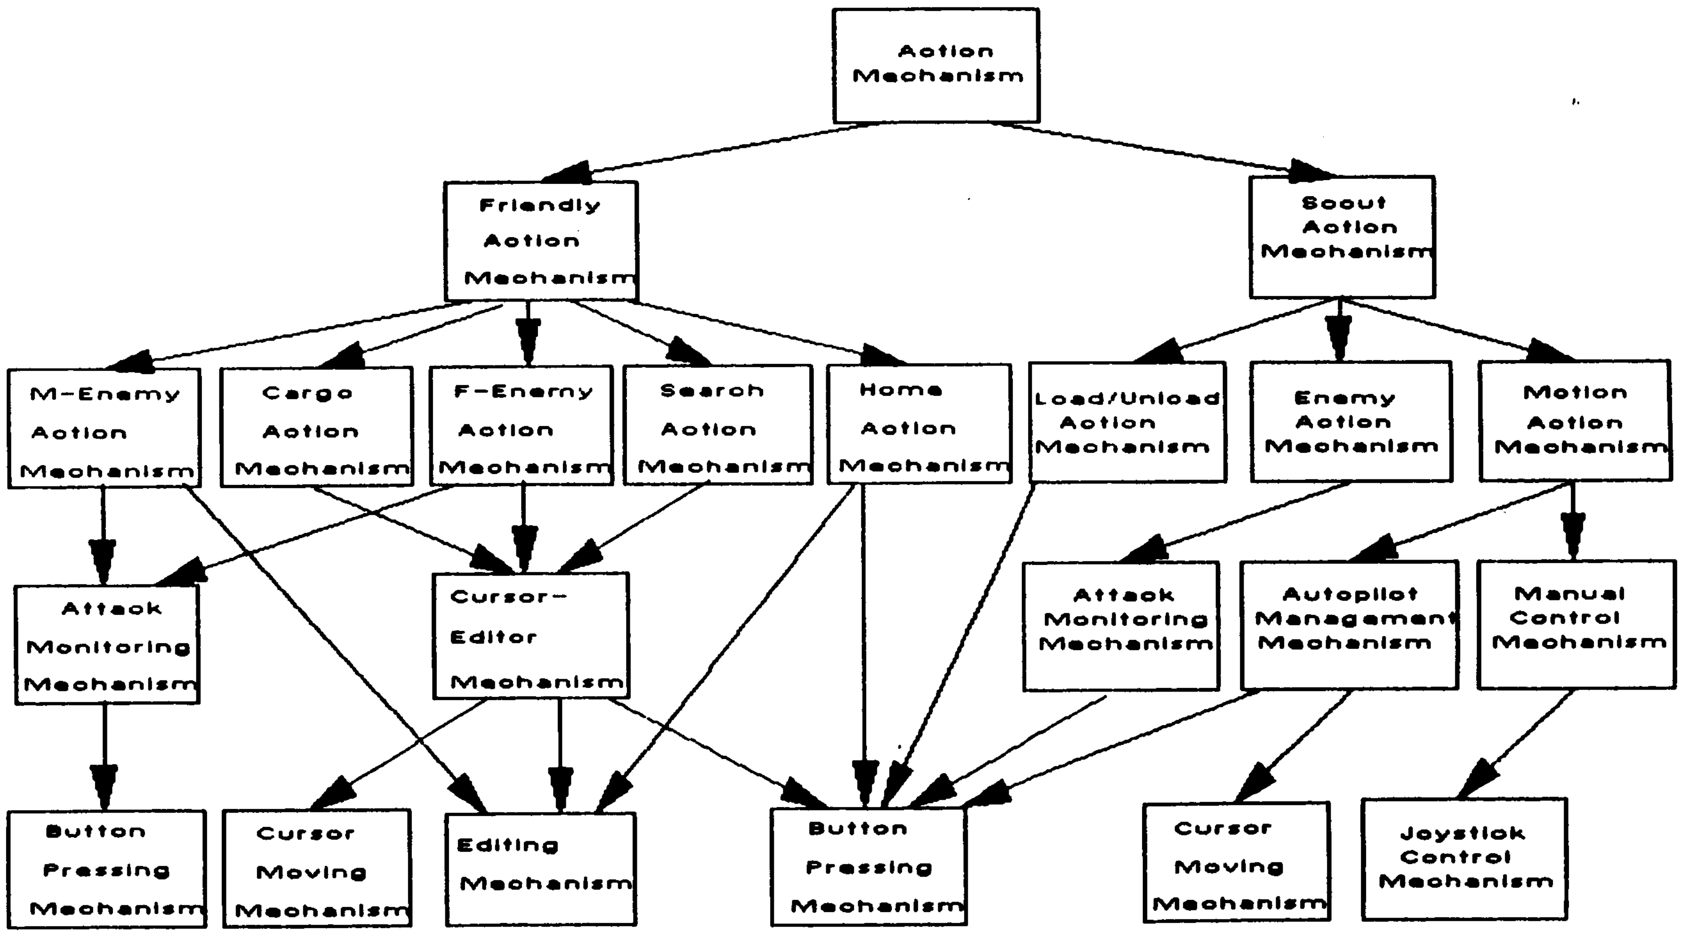

Note: The button and cursor mechanisms have been replicated at the bottom level of the diagram for clarity

Action mechanisms

Figure 16

of the action mechanisms in the present model can be viewed as a similar type of construct.

The links between the mechanisms denote information flow. In the diagram, the direction of information flow is always downward. For example, the most global action mechanism (at the top of the diagram) is activated each time any interface activity is required. It then determines whether the need for action concerns the friendly craft or scout action mechanisms, and sends information downward to the appropriate mechanism. Although information always flows downward in this component of the model, an implicit upward flow of information does exist through the external environment. That is, the results of the actions performed by the lower level mechanisms can become known to the upper level mechanisms through the perception of the environmental changes caused by those actions.

The most general action mechanism at the top of the diagram is responsible for scheduling the possibly competing demands for scout-related and friendly-related interface control activity. To perform this function, the action mechanism operates upon the output of the action selection mechanisms. This output is a matrix indicating the desired activity for the scout and each of the friendly craft. The action mechanism is sensitive to changes in this matrix. A change in the currently desired activity for the scout or a friendly craft indicates that interface activity is required to bring that craft's activity into conformity with its desired activity. Each iteration, all such changes are noted along with their criticality. In the present model, the criticality of a potential change of activity directly corresponds to the affordance level of the desired action. It might, though, be more realistic to assume that the criticality of a potential change of activity should be a function of the difference between the affordance levels of the current and desired actions, although the model has not been constructed according to this assumption. This assumption would probably better capture the notion that the most important actions are those that resolve the greatest difference between the current and desired world state. In the present model, though, the most important actions are defined to be those that are associated with the greatest affordance level.

After the most critical change in activity is identified, either the friendly craft action or scout action mechanism is activated to perform the interface activities required to implement the change. The friendly craft action mechanisms appear on the left side of the diagram and the scout action mechanisms appear on the right side. The highest level mechanisms on each of the two sides are responsible for selecting the lower level mechanisms that are appropriate for performing the interface activity. For the friendly craft, these are the mobile enemy action, cargo action, fixed enemy action, search action, and home action mechanisms. For the scout, these are the load/unload action, enemy action, and search action mechanisms.

These action mechanisms are directly responsible for coordinating activity that requires interaction with multiple interface controls. The interface controls are: the friendly craft text editor, used to enter action commands for the friendly craft; the world display cursor used to enter search waypoints; the set of pushbuttons used to activate the editor, to provide real-time modification of friendly action commands, and to control the scout auto-pilot; and the scout control joystick, used for manual control of scout motion.

As discussed in Chapter III, the crews differed in the way in which they interacted with the interface controls, so some of the action mechanisms had to be individually tailored to describe each crew. Crew 1, a one-person crew, used the autopilot for control of the scout, presumably in order to cope with the extensive task demands associated with having to control both the scout and friendly craft. Crew E, the expert one-person crew, used manual control of the scout that was intermittently interrupted by editing sessions for friendly craft control. For the two-person crew, Crew 2, the pilot was dedicated to the task of scout manual control, while the navigator was dedicated to the task of friendly craft control. Since some of the action mechanisms in the model are dependent on the crew being described, a detailed description of these mechanisms will be given in the following chapter which describes how the model was parameterized in an attempt to mimic the behavior of each of the human crews.

CHAPTER VI

FITTING THE MODEL TO HUMAN CREWS

Introduction

This chapter describes how the model was parameterized to mimic the behavior of the three human crews. First, the method used to select model parameters is discussed. This method is guided by both theoretical and empirical considerations. The theoretical considerations refer to identifying those manipulations of the model that have meaningful psychological interpretations. The empirical considerations refer to the need to demonstrate the model's sufficiency through adequately matching crew behavior. Next, a set of crew-specific parameter values are described that were used to generate model behavior similar to the behavior exhibited by each of the crews. In some cases the parameters are numerical values, but in other cases the value of a parameter indicates a particular policy, such as the use of the autopilot for scout motion control. Naturally, the minimal set of model parameters were sought that could be adjusted to match empirical data. Finally, the method used to assess the model's empirical adequacy is described. The results of the crew modeling are discussed in the following chapter.

Parameterization Approach

Due to the complexity of the model, some rationale is needed to guide the identification of the parameters that can be manipulated to match crew performance. Any division of the model into constants, parameters, and variables must be based on implicit (at least) assumptions about the properties of the psychological mechanisms to which the model's constructs refer. One desire is to account for the behavioral differences between the three crews in the most parsimonious fashion. At a minimum, the characterization of crew differences in terms of the model's constructs should be simpler than the characterization of crew differences in terms of the empirical data. Consider, for example, the differences in the sets of condition-action rules required to differentiate two human subjects in a production system model of problem solving. If these differences cannot be demonstrated to be a simpler description of inter-subject differences than a description of inter-subject differences in purely empirical terms, there are serious questions as to the theoretical contribution of the production system model.

Thus, one test of the present model is the degree to which it can economically account for crew behavioral differences in ways that are readily interpretable in terms of the model's constructs. At the most general level, the model can be decomposed into perceptual, selection, and action components. Each of these components can be further decomposed into a set of specific mechanisms, as described in the previous chapter. The primary goal of this chapter is to explain how and why particular mechanisms were chosen to describe each of the three human crews.

Most of the emphasis has been placed on the development of the perceptual components in the model. This is due to the fact that the modeling approach is based on viewing skilled human performance as relying heavily on powerful, task-dependent, context-specific perceptual mechanisms. The identification and development of the selection, or central, components of the model was secondary. This task was guided by considering what information-processing tasks could not be easily accounted for by the perceptual mechanisms. In this way, the role of the selection mechanisms can be seen to be implicitly defined by the specification of the perceptual mechanisms. Finally, consideration was given to the action mechanisms that were responsible for simulating the observable crew activities. Little theoretical consideration has been given to the design of the action mechanisms other than the desire to employ a hierarchical structure with enough flexibility to account for the crews' different physical interface control behaviors.

Given this overview of the modeling emphasis, the following approach for parameterizing the model was used. First, action mechanisms for each crew were developed in accordance with the observable interface control policies used by the three crews. Mean editing times and policies for scout and friendly craft control were identified from the empirical data. In the case of the scout, the relevant data concerned whether manual or autopilot control was used. For the friendly craft, the action commands entered via the text editor that were used to control the friendly craft were identified. Videotapes of crew sessions were consulted to measure average times to press buttons and manipulate the joystick controlling the cursor on the world display.

Given these easily measured differences between crews, perhaps the simplest hypothesis to be considered is that the array of performance differences identified in the crew profiles could be due entirely to these parameters. This would be an attractive hypothesis from a methodological perspective since it would confine the parameterization search to the action component of the model, allowing the perceptual and selection components to remain invariant over different crews. An analysis of the empirical data, though, proves this hypothesis untenable. While different action mechanisms might account for certain performance differences, such as craft idle time for example, an examination of the entire crew profiles indicates that the crews also differed in terms of how well they made use of the five craft in ways that were independent of the efficiency by which the crews interacted with the physical interface.

For example, Crew 1, the one-person crew who used autopilot control of the scout, appeared to have used a inefficient policy to determine scout activities. As described in the crew performance profiles in Chapter III, this crew's inability to search the world with the scout did not seem to be entirely due to the scout's decreased speed and higher fuel usage under autopilot control. Rather, searching inefficiency also resulted from the fact that Crew 1 used the scout to load more cargo than the other one-person crew, thereby keeping the scout from searching the world for other cargo and enemy craft.

Another example of crew differences that do not appear to be the result of differences in interface manipulation skills concerns Crews 2 and E. Although Crew E was less able to keep the scout in motion than the two-person crew, this subject was able to cover more unsearched area with the scout than Crew 2. This fact suggests that Crew E's scout search paths were somehow more efficient, although Crew 2 was able to search the same total area as Crew E due to searching performed by the friendly craft. Another substantive difference between these two crews was that Crew E was able to unload a higher percentage of loaded cargo than Crew 2. This appears to reflect a difference in the degree to which the friendly craft could be effectively managed, and seems difficult to explain in terms of differences in interface manipulation skills.

These crew differences have been discussed because they appear to provide evidence that the crews differed in ways in addition to any differences that were observed in terms of interface manipulation skills. The data cited above seem to indicate the presence of crew differences at a strategic level, rather than as variations in the efficiency with which crews could interact with the physical interface. It could be the case, though, that "low-level" differences in interface manipulation skills may be manifested, albeit in rather complex ways, in terms of "high-level" strategic differences. That is, it might be the case that high-level strategic differences are emergent upon differences in terms of low-level interface manipulation skills.

For example, the strategic differences exhibited between two football teams (e.g., a run-oriented versus a pass-oriented offense) might be due in large part to differences between the two teams in terms of the physical attributes and execution skills possessed by the teams' players. The two different strategies might be the result of the same type of strategy selection process simply operating upon different sets of information. One such process might attempt to determine an optimal offensive strategy as a function of the execution skills of the team's players. Given teams composed of players with different sets of skills, different optimal offensive strategies might be produced by this process.

It may be the case that this phenomenon could help explain the strategic differences observed in the laboratory task. If this is the case, it would appear possible to construct a model capable of determining an appropriate strategy based upon information pertaining to interface manipulation skills. The different strategies produced by such a model might be capable, then, of accounting for the observed strategic differences between the crews in terms of differences solely at the level of interface manipulation skills. If this approach were to be taken in the current work, the parameters of the perceptual and cognitive mechanisms might not need to be altered to generate different behaviors at the strategic level.

As a matter of fact, this approach has been used to describe strategic differences between crews in terms of autopilot usage (Kirlik, 1987). A model was constructed to demonstrate that different strategies for autopilot usage (e.g., dedicated autopilot usage, dedicated manual control, switching between the autopilot and manual control based upon editing demands) could be shown to be the result of an optimization process operating upon data concerning a crew's ability to perform manual control. That is, it was shown that some of the different autopilot usage strategies observed in the experiments were consistent with the assumption that crews were attempting to optimize average scout speed given their differing ability to perform manual control. Crews highly skilled at manual control, for example, could be shown to be acting optimally by not using the autopilot. Crews with low manual control skills could be similarly shown to be acting optimally by using dedicated autopilot control. In addition, it was shown that the strategy of engaging autopilot control before editing and returning to manual control after editing was never optimal regardless of the crew's manual control proficiency. This strategy was always sub-optimal due to the large amount of time required to set-up and engage the autopilot. Thus, the model provided one explanation of the fact that no crews were observed to use this strategy. In this model, then, strategic differences were shown to be emergent upon differences in terms of manipulation skills.

It might, then, be possible to construct a similar model for showing how the strategic differences between the crews could be the result of differences in terms of low-level interface manipulation skills. Such a model, though, would not be consistent with the assumptions discussed in Chapter IV concerning the nature of the real-time operation of the mechanisms underlying skilled performance. It was argued that the operation of these mechanism is ahistorical and non-teleological. In other words, it was suggested that a description of the real-time operation of these mechanisms is independent of the previous shaping forces that are sensitive to task goals and result in expert or possibly optimal performance. Thus, an appropriate description of the operation of these mechanisms need not refer to the past events that were responsible for their existence or the goals to which their operation contributes. It is assumed that optimization operations are not a component of the real-time operation of the mechanisms underlying skilled performance, although optimization (or learning) operations might be responsible for the "design" of the mechanisms at any point in time.

It is therefore suggested that a model which demonstrates how high-level strategic differences would be the (perhaps optimal) emergent results of low-level, skill-related differences is not appropriate for describing the real-time operation of the mechanisms underlying skilled performance. While the autopilot usage model discussed above, for example, might be useful for predicting which mechanisms will be evolved as a result of the strategy selection process (i.e., those that implement the selected strategy), it would not be useful for describing the real-time operation of these mechanisms. To use the model for this purpose would amount to making the dubious claim that optimization operations were performed each time the crew took an action related to the use of the autopilot or the manual control joystick. In direct contrast to this claim, the present modeling approach assumes that the strategy selection operations that might contribute to highly skilled performance are performed off-line, or incrementally, and should be considered distinct from the mechanisms that actually implement the selected strategy. To return to the football case, for example, the process which operates upon information concerning the players' execution skills to determine an offensive strategy would probably be performed by the coach prior to the game, or in an incremental fashion while watching the game from the sidelines. A real-time description of the behavior of the players while implementing this strategy need not be concerned with the coach's strategy selection process.

While it might be possible to describe crew strategic differences as emergent upon differences in interface manipulation skills, then, this description would exist at a different level than a description of strategic differences in terms of differences within the real-time mechanisms underlying performance. The first description would provide an account of why the mechanisms were "designed" in a certain way, and would probably make reference to task goals and the task feedback that would be relevant to obtaining these goals. This level of description might, then, be defined as teleological in nature. The second description, on the other hand, would provide an account of the operation of the mechanisms and would not be concerned with task goals or learning. This description might therefore be defined as non-teleological or mechanistic. Since this research is devoted to describing the real-time operation of the mechanisms underlying skilled performance, only the second type of description will be used here, although pursuing the "design" issues would surely provide a more comprehensive account of the nature of skilled performance.

A direct implication of the decision to ignore the teleological level of description is that the process of fitting the model to crew behavior might appear shallow or descriptive in nature. This issue will be discussed in greater detail in the next sections where the process of model parameterization is described. At this point, though, it should be clear that the adoption of a mechanistic level of description entails that a description of the strategic differences between crews will only concern hypothesized differences in the real-time mechanisms that implement these strategic differences. Based on the arguments presented above, it does not appear possible to completely account for strategic differences in terms of action mechanisms alone, since these mechanisms are only responsible for implementing, rather than determining, a particular strategy.

Given that it does not appear possible to completely account for crew differences in terms of action mechanisms alone, the next simplest hypothesis to be examined is that crew differences were also due to the existence of different perceptual mechanisms, but crews did not differ in terms of the selection, or central, mechanisms. The hypothesis that the crews' perceptual mechanisms differed would appear to be consistent with the theoretical ideas on which the modeling approach was based. This hypothesis suggests that crews differed in terms of the degree to which their perceptual mechanisms became attuned to the affordances in the simulated world that would allow for rapid and effective action selection.

No justification, on the other hand, is given in the theoretical arguments for the assumption that the selection mechanisms were invariant over the three crews. Rather, this hypothesis is adopted primarily for methodological reasons. If the model can be made to mimic the different behaviors of the human crews without manipulating the central selection mechanisms, little empirical justification can be provided for the hypothesis that the crews' selection mechanisms differed. On the other hand, if the model can be made to mimic crew differences while holding the perceptual mechanisms fixed and varying the selection mechanisms, little empirical justification could be provided for the hypothesis that the crews' perceptual mechanisms differed. If neither the perceptual nor selection manipulations alone is capable of mimicking crew differences, then the hypothesis that both these mechanisms differed between crews would appear to be required.

The current approach for describing crew differences is based on the requirement that simpler descriptions should be explored before more complex descriptions are considered. Therefore, the hypothesis that a description of crew differences can be localized to either the perceptual or selection mechanisms, but not both, will be explored before considering the hypothesis that crews differed in regard to both of these mechanisms. For reasons discussed below, it was decided that the selection mechanisms should be held constant and the perceptual mechanisms should be manipulated in the attempt to describe crew differences.

One reason that the selection mechanisms were held constant to describe crew differences is due to the limited ability to vary these mechanisms in theoretically meaningful and empirically relevant ways. The operation of the three selection mechanisms in the model, the scout waypoint selection mechanism, the friendly waypoint selection mechanism, and the object selection mechanism can be characterized in the same simple way. Each of these mechanisms selects the action with the maximum affordance value subject to a small set of constraints. The object selection mechanism is more complex since it is capable of predicting future world states to enable planning. The constraints used in each of the mechanisms are specific to the action selection task. For example, the scout waypoint selection mechanism uses a constraint to reduce backtracking, and the object selection mechanism uses a constraint to eliminate competition between the five craft.

Given this characterization of the selection mechanisms, there would appear to be a very small set of possible manipulations that could be used to fit the behavior of the different crews. The three most obvious manipulations would be the following: a) altering the assumption that crews were choosing the actions with maximum affordance value; b) altering the assumed set of constraints used by the selection mechanisms; and c) altering the assumptions concerning the planning horizon used to select future craft actions. Unfortunately, manipulation (a) seems to be of little theoretical utility and manipulations (b) and (c) do not seem to be empirically relevant.

Abandoning the assumption that crews were selecting actions with maximum affordance value (manipulation (a)) would appear to cause two theoretical difficulties. The first would be concerned with the problem of identifying a coherent set alternative assumptions that could be used for manipulating the operation of the selection mechanisms. Do crews choose the actions with the second highest affordance value, or perhaps actions with affordance values in a specified range? These assumptions would appear to be arbitrary and of little theoretical value. The second problem with this approach is that the objection could always be made that crews were indeed choosing the maximally afforded actions, but the reason that a more complex selection operation was required to mimic behavior was that the affordance values were simply generated in the wrong fashion. That is, it would appear at least theoretically possible to mimic the operation of a more complex selection mechanism with a simple maximization-based mechanism by changing the way in which the affordance values are calculated. This objection would essentially claim that crews were indeed maximizing but they were maximizing with respect to a different set of criteria than used by the more complex, non-maximizing, mechanism. In the terminology of the present model, then, the argument could always be made that the crews' actually differed with respect to their perceptual mechanisms (which generate affordance values) rather than with respect to their selection mechanisms. Therefore, the use of manipulation (a) for varying the selection mechanisms appears to be of limited theoretical utility.

Similarly, it does not appear theoretically rewarding to vary the constraints used by the selection mechanisms in an attempt to describe crew differences. The role played by these constraints is to eliminate certain highly afforded actions from consideration for reasons specific to the action selection task. For example, the exclusivity constraint used by the object selection mechanism eliminates candidate craft actions that compete with actions to be taken by another craft. In the scout waypoint selection mechanism, constraints are used to eliminate candidate waypoints that would cause the scout to run out of fuel or perform extensive backtracking. The observed behavior of each of the three crews appeared to be consistent with these constraints.

The reason it does not appear useful to vary these constraints to describe crew differences is quite similar to the reason that manipulations of the maximization assumption were rejected. As was the case with varying the maximization assumption, it appears to be the case that the need to adjust the constraints used by the selection mechanisms could be eliminated by suitably chosen manipulations of the perceptual affordance mechanisms. The affordances provided by the perceptual mechanisms generate a candidate set of actions which are then pruned by rejecting those actions that do not satisfy the constraints. Since the constraints operate upon the affordances in this way, it would appear possible that the effects of introducing a new constraint could be mimicked by a suitable readjustment of the affordance value calculations. Once again, then, the argument could always be made that the selection mechanisms were indeed invariant over crews and that crew behavioral differences were due solely to differences in their perceptual mechanisms.

These arguments concerning the possibility of behaviorally indistinguishable tradeoffs between the operation of the perceptual and selection mechanisms raise the possibility that the design of the model is underspecified by observable behavior. For example, it has been suggested above that a desired change in the behavior of the model could be accomplished by manipulations of either the constraints used in the selection mechanisms or the parameters used in the perceptual affordance value calculations. In model construction, then, choices had to be made to determine which operations would be located in the perceptual mechanisms and which would be located in the selection mechanisms. The assignment was not arbitrary. Rather, constraints were included in the selection mechanisms for only those operations for which a perceptually-oriented processing mechanism could not be identified. A similar type of tradeoff was noted between the maximization operation and the operation of the perceptual affordance mechanisms. The existence of these potential tradeoffs would cause a serious problem if the model was being used in an attempt to prove whether certain operations were being performed perceptually or centrally. These tradeoffs are not a problem in this work, though, since the goal is merely to demonstrate that the behavior of a strongly perceptually-oriented model is consistent with observable behavior.

Since it does not appear to be theoretically rewarding to vary the maximization operation or the constraints used by the selection mechanisms, manipulation (c) seems to be the only remaining candidate for describing how crews' selection mechanisms may have differed. Unlike manipulations (a) and (b), manipulation (c) appears to be theoretically interesting since it would claim that crews differed in terms of the planning horizon they used to select future craft actions. This manipulation, though, does not seem to be able to produce variations in model behavior that were consistent with the observed variations in crew behavior.

As described in the previous chapter, the design of the object selection mechanism was based on the assumption that crews could plan three actions into the future for each craft. This assumption was based on an informal sensitivity analysis of the model's behavior as a function of planning horizon. This informal analysis was performed before the model was fully parameterized to mimic the behavior of each of the crews. To achieve reasonable model behavior, this analysis indicated that it was critical to assume a planning horizon of at least two actions, it was less critical but still necessary to assume a planning horizon of three actions, but increasing the planning horizon beyond three actions was of limited additional benefit. Since it was not desirable to assume a planning process of greater complexity than was necessary, a planning horizon of three actions was assumed.

To provide a more formal justification for the need to assume even this level of model complexity, a similar sensitivity analysis was performed after the model was fitted to crew behavior. After model parameters were chosen to provide the best fit to Crew E's behavior (see the following sections), the model's planning horizon was then reduced to one action. The behavior of the model with the one-action planning horizon (the short-term model) was then compared with the behavior of the model with the three-action planning horizon (the long-term model) to identify how the reduction in planning horizon affected the behavior of the model. As described below, the nature of the sub-optimal behaviors generated by the short-term model did not appear to be consistent with the behavior of any of the three crews. This result provides post-hoc evidence for the necessity of assuming the more complex, three-action planning horizon.

When compared with the long-term model used to mimic the behavior of Crew E, the performance of the short-term model was inferior in terms of the average number of points scored per mission. (See the last section of this chapter for a description of the model comparison process). The average points scored by the short-term model did not differ significantly from the average points scored by Crew 1. This finding suggests that it might be possible to account for Crew 1's inferior performance in terms of the assumption of a reduced planning horizon. A detailed analysis of the behavior of the short-term model, though, indicates that the sub-optimal behaviors that contributed to its inferior performance were not consistent with Crew 1's behavior. As will be discussed below, the assumption of a short-term planning horizon produces sub-optimal behaviors that did not appear to be consistent with the behavior of any of the three crews.

In terms of the set of performance measures used to describe crew behavior, the major factor that contributed to the short-term model's poor performance was an inability to score points by unloading discovered cargo at home base. The short-term model and Crew 1 were not significantly different in terms of the number of cargo unloaded per session, (and both unloaded less cargo than did Crew E), but the factors that contributed to their similarly poor performance on this measure were found to be quite different. As discussed in Chapter III, Crews 1 and E were similar in terms of the percentage of discovered cargo that were eventually unloaded at home base, but Crew 1 discovered and returned less cargo than did Crew E. That is, these two crews were similarly efficient in unloading the cargo that had been discovered, but Crew E discovered more cargo than Crew 1 and was thus able to unload more cargo to score points.

In contrast to this characterization of the differences between Crews 1 and E in terms of cargo processing, the difference between the short-term model and Crew E concerned the percentage of discovered cargo that were unloaded at home base, rather than the number of cargo that were discovered. The short-term model and Crew E did not differ significantly in terms of the number of cargo that were discovered, (and they both discovered more than Crew 1), but the short-term model was able to unload a smaller percentage of these cargo at home base. In fact, the short-term model was significantly worse on this measure than any of the three crews, as it was only able to unload an average of 61% of discovered cargo at home base. (No crew unloaded less than 80% of the discovered cargo, on average, per mission). Therefore, a major factor that contributed to the short-term model's sub-optimal performance was an inability to unload the cargo that had been discovered, whereas a major factor that contributed to Crew 1's sub-optimal performance was an inability to discover cargo.

The short-term model, then, exhibited a pronounced inability to efficiently process discovered cargo. This sub-optimal behavior was not exhibited by any of the three crews. The short term model's sub-optimal behavior in this regard can be explained by examining the role of the planning horizon in generating model behavior. It should be pointed out that, in this work, the term "planning horizon" has a technical meaning that is specific to the structure and behavior of the model. The planning mechanism used in the model is event-based, rather than time-based, and therefore the length of the planning horizon indicates only the number of actions into the future that the model is able to consider. Since these actions have variable duration, the length of the planning horizon does not have a tight correspondence with the amount of time into the future that the model can plan ahead.

The assumption of a one-action planning horizon constrains the action selection mechanism of the model to the task of determining only the most appropriate <u>current</u> action for each craft. Therefore, the competition between craft for consumable objects, such as cargo, is concerned only with each craft's current affordance value for the object, and does not consider the affordance value that a craft would have for the object at the termination of its current action. This restriction, on its own, can be seen to be sufficient for producing the short-term model's sub-optimal cargo processing behavior.

To see how the inability to efficiently unload discovered cargo can arise from restricting the planning horizon to a single action, consider the case where a craft is currently traveling to load a piece of cargo or to return to home base. After the craft has been committed to this action, assume that a piece of cargo is discovered nearby the destination of the craft. Assume, also, that the affordance value for loading this new cargo is not great enough to exceed the craft's affordance value for the current action so that the craft's current action is not interrupted. In this case, the craft is not considered as a candidate for loading the new piece of cargo since it has a more highly afforded current action. Even though, for example, the craft would have a very high affordance for loading the new piece of cargo at the termination of its current action, this fact is not relevant in determining which craft will be allocated to the cargo. The cargo will then be allocated to a second craft, possibly many miles away, that has the highest current affordance value for the new piece of cargo. This craft may not even be able to reach the cargo until well after the original craft could have loaded the cargo by proceeding to its location after the termination of its current action.

Therefore, the most appropriate action of having the original craft load the cargo after completing its current action has been neglected in favor of using a craft that may take a much longer time to accomplish the task. Given the way in which the action selection mechanism operates, the original craft does not compete for the new cargo even after the termination of its current action since the cargo has already been allocated to a different craft. Therefore, the original craft may be given a search action at the completion of its current action even though it still may be closer to the cargo than the second craft to which the cargo has been assigned. Perhaps, though, the selection mechanism could be improved by allowing the original craft to compete for the cargo at the termination its current activity, thereby possibly interrupting the activity of the second craft. Even with this adjustment, though, the resulting behavior would still be inefficient due to the amount of wasted time spent by the second craft traveling to a cargo that it will never eventually load.

Although the conditions that give rise to this cargo processing inefficiency might appear too complex and unlikely, in actuality this phenomenon was observed to occur with reasonable frequency. These conditions occurred with at least enough frequency to cause the short-term model to unload a significantly lower percentage of discovered cargo than any of the three crews. For example, there were many cases in which a friendly craft discovered a cargo on its way to home base in urgent need of fuel or missiles. In the cases where this cargo was itself nearby home base and the other friendly craft were distant from home base, the most appropriate plan of action would be to have the craft that discovered the cargo return to load it after completing its trip home. Due to the operation of the craft allocation policy described above, a craft currently searching at a large distance from home base would be assigned to load this cargo. As a result of this type of inefficiency, the missions performed by the short-term model typically ended with a large number of discovered cargo not returned to home base.

None of the crews were observed to produce sub-optimal behavior that was consistent with the assumption of a restricted planning horizon. For this reason, it was decided that the planning horizon would not be manipulated in an attempt to mimic crew behavior. In summary, then, each of the three possible selection mechanism manipulations have been examined and found lacking for either theoretical or empirical reasons. Therefore, the selection mechanisms were assumed invariant over each of the three crews.

Recall that the model contains three selection mechanisms: the scout waypoint selection, friendly waypoint selection, and craft object selection mechanisms. It should be emphasized that the invariance assumption for these mechanisms does not imply that the same scout paths, friendly waypoints, and craft object allocations were assumed to be generated by the different crews. Rather, the selection mechanisms have a relatively minor role in action generation. Their function is to select in a simple way from the candidate set of action alternatives provided by the perceptual mechanisms. In the case of friendly craft waypoint selection, the action selection mechanism operates by finding the action alternative (a search waypoint) with the greatest affordance level. The scout waypoint and craft object selection mechanisms also operate by selecting action alternatives with the greatest affordance level, subject to appropriate constraints concerning backtracking avoidance in the case of the scout, and mutual exclusivity in the case of the friendly craft.

Given this functional characterization of these selection mechanisms, it is easy to see that the major factor determining which actions get selected is the perceptual component of the model that defines the action alternatives and generates the dynamic, context sensitive, affordance level for each of the alternatives. Therefore, once the action mechanisms were constructed for each crew and a single set of selection mechanisms were created, the model parameterization exercise focused on designing perceptual mechanisms that would provide the affordance structures that were consistent with the behavior of each of the three human crews.

Detailed descriptions of each of the action and perceptual mechanisms used to mimic the behavior of each of the crews are given in the following sections. The action mechanisms were designed to be consistent with the empirical data, and were left unchanged during the model fitting process. Initial estimates for the parameters of the perceptual mechanisms were generated based on the crew performance profiles. These parameters were subsequently tuned to gain better matches with the empirical profiles in an iterative analysis of model behavior as a function of the parameters of the perceptual mechanisms.

Action Mechanism Parameters

The action mechanisms shown in Figure 16 in the previous chapter were parameterized for each crew by referring to empirical data concerning crew physical activities. Videotapes were consulted to measure the timing parameters used in the lowest level mechanisms that simulated observable editing, button-pushing, and joystick manipulation behavior. These parameters indicated the total time spent in both the preparation and execution of these actions so that the model would be diverted from engaging in other tasks for the same duration as were the human crews.

Table 2 on the following page describes the three highest level mechanisms: the Action Mechanism, the Friendly Action Mechanism, and the Scout Action Mechanism. As indicated in the table, the only difference between the three crew models with respect to these three mechanisms concerns whether or not friendly- and scout-related interface control activity could be performed serially or in parallel. The two models for describing one-person crews, Crews 1 and E, are constrained to perform friendly and scout activities serially. The model used to mimic the behavior of the two-person crew, Crew 2, can perform friendly and scout activities in parallel. For example, the Crew E model interleaves scout manual control with entering friendly craft commands in the text editor, while the Crew 2 model performs these two tasks simultaneously.

TABLE 2

HIGHEST LEVEL ACTION MECHANISMS

<u>Mechanism Name</u>	:	Action Mechanism
Description	:	Identifies requests for interface activity from affordance distributions and activates either the scout or friendly craft action mechanisms.
Parameters	:	Serial vs. parallel interface control for scout and friendly craft activities
Parameter Values	:	Crew 1 Model - Serial
	:	Crew E Model - Serial
	:	Crew 2 Model - Parallel

<u>Mechanism Name</u>	:	Friendly Action Mechanism
Description	:	Identifies whether friendly activity request pertains to mobile enemy craft, fixed enemy craft, cargo, search, or home activity, and activates the appropriate mechanism.
Parameters	:	None

<u>Mechanism Name</u>	:	Scout Action Mechanism
Description	:	Identifies whether scout activity request pertains to enemy craft, loading/unloading, or scout motion, and activates the appropriate mechanism.
Parameters	:	None

Table 3 on the following pages describes the seven friendly craft action mechanisms that coordinate the simulated physical interface activities associated with friendly craft control. The parameter values used in the first five of these action mechanisms indicate the action commands used to assign activities to the friendly craft. These action commands were identified directly from the empirical data.

The Mobile Enemy Action Mechanism and Cargo Action Mechanism did not differ in the three crew models since all three crews used the same action command to attack mobile enemy craft and to command friendly craft to load cargo. Similarly, Crews 2 and E used the same action command to attack fixed enemy emplacements, so the Fixed Enemy Action Mechanisms did not differ between these two crew models. Crew 1 never attacked fixed enemy emplacements with the friendly craft, so this mechanism was never activated in the model used to describe this crew.

In the Search Action Mechanism, the behavior of Crews 1 and 2 was approximated by using the friendly craft search command "S:^,G", while the behavior of Crew E was approximated with the search command "S:^,G,P". Each crew actually used a range of different search commands. For example, all three crews occasionally used a search command with multiple goto waypoints (e.g., S:^,G,G), although such commands were very rare for Crews 2 and E. In addition, all

TABLE 3

FRIENDLY ACTION MECHANISMS

<u>Mechanism Name</u>	:	Mobile Enemy Action Mechanism
Description	:	When activated by the Friendly Action Mechanism, activates the Editing Mechanism, inputs the attack action command, then activates the Attack Monitoring Action Mechanism.
Parameters	:	Attack action command.
Values	:	All crew models - E:A,A,A
<u>Mechanism Name</u>	:	Cargo Action Mechanism
Description	:	When activated by the Friendly Action Mechanism, activates the Cursor - Editor Mechanism to move the world display map cursor to the cargo location and to enter the text editor, inputs the cargo action command.
Parameters	:	Cargo action command
Values	:	All crew models - C:^,G,v,L,^
<u>Mechanism Name</u>	:	Fixed Enemy Action Mechanism
Description	:	When activated by the Friendly Action Mechanism, activates the Cursor - Editor Mechanism to move the world display map cursor to the enemy location and to enter the text editor, inputs the fixed enemy attack action command.
Parameters	:	Fixed Enemy action command
Values	:	Crew 1 Model - None
	:	Crew E Model - E:^,G,A,A,A
	:	Crew 2 Model - E:^,G,A,A,A

TABLE 3
(continued)

FRIENDLY ACTION MECHANISMS

<u>Mechanism Name</u>	:	Search Action Mechanism
Description	:	When activated by the Friendly Action Mechanism, activates the Cursor - Editor Mechanism to move the world display map cursor to the search waypoint location and to enter the text editor, inputs the search action command.
Parameters	:	Search action command
Values	:	Crew 1 Model - S:^,G
	:	Crew E Model - S:^,G,P
	:	Crew 2 Model - S:^,G
<u>Mechanism Name</u>	:	Home Action Mechanism
Description	:	When activated by the Friendly Action Mechanism, activates either the Cursor - Editor Mechanism or the Editing Mechanism to move the world display map cursor to home base (if necessary) and to enter the text editor, inputs the home action command.
Parameters	:	Home action command
Values	:	Crew 1 Model - X:U (no cursor required)
	:	Crew E Model - X:U (no cursor required)
	:	Crew 2 Model - S:^,G,v,U (requires cursor)
<u>Mechanism Name</u>	:	Attack Monitoring Mechanism
Description	:	When activated by either the Mobile Enemy Action Mechanism or Fixed Enemy Action Mechanism, activates the Button Pressing Mechanism to enable the firing of missiles until either the enemy or friendly craft is destroyed.
Parameters	:	None

TABLE 3
(continued)

FRIENDLY ACTION MECHANISMS

Mechanism Name	:	**Cursor - Editing Mechanism**
Description	:	**When activated by any mechanisms that require cursor manipulation followed by editing, activates the Cursor Mechanism and Editing Mechanism in sequence.**
Parameters	:	**None**

three crews occasionally dispensed with the up (^) action, although Crew E did this very rarely. Unfortunately, no consistent policies for search command selection could be identified. For example, it was hypothesized that, for Crews 1 and 2, whether or not the up (^) action was used was a function of the current craft altitude. This hypothesis could be rejected due to the fact that up actions were used inconsistently at the beginning of the mission when all craft were at the same altitude. The hypothesis that the selection of the up action was due to planned travel through forested locations received some support, at least for Crew 1, although this policy was not consistently applied.

Due to this inability to identify consistent policies for search command selection, the simple search commands described above were used. The only difference between the search commands for the three crews pertains to the use of the patrol action. Crews 1 and 2 never used the patrol (P) action, while Crew E nearly always terminated search action commands with patrol actions. The patrol action keeps the friendly craft searching in a circular pattern after the termination of the goto action. The search commands used to describe the behavior of Crews 1 and 2 did not include patrol actions, resulting in the friendly craft sitting idle at the termination of search commands.

The Home Action Mechanism also had to be individually tailored to describe the behavior of the three crews. Crew

E consistently used the "X:U" command, whereas Crew 2 consistently used the "S:^,G,v,U" command. Both commands send the friendly craft home to unload, although the second command allows the operator to specify the exact location (within 0.4 miles of home base) at which the craft should stop and unload, whereas the first command always sends the friendly craft to the exact center of home base. The second command was presumably adopted to lessen the potential for collisions between friendly craft that would be caused by commanding multiple craft to home base simultaneously. The first command, though, has the benefit that the cursor does not have to be manipulated to enter a home command.

For describing Crew 2, then, the home command requiring cursor manipulation was used in order to ensure that the model and human crew performed interface manipulation activity of the same duration when commanding a friendly craft to go to home base. This model does not, though, attempt to scatter the destination points of multiple friendly craft around home base to avoid collisions. (The reason that the model does not perform this function is discussed at the end of this chapter). For describing Crews 1 and E, the model uses the home command that does not require cursor manipulation. As was the case with search commands, the use of a single home command for describing Crew 1 is a simplification, since this crew occasionally used the command employed by Crew 2.

The action mechanisms used for organizing the simulated physical activity associated with control of the scout are described in Table 4 on the following pages. The only mechanism with parameter values that differ between the three crew models is the Motion Action Mechanism. The models used to describe Crews 2 and E simulate the activity associated with manual control of the scout, whereas the model used to describe Crew 1 simulates the scout autopilot management physical activity. The approximation of dedicated autopilot control was used, although Crew 1 did engage manual control briefly and intermittently.

The Manual Control Mechanism, used in the models of Crews 2 and E, is activated by the Scout Action Mechanism to generate a path to a desired scout waypoint. The desired waypoint, provided by either the Scout Waypoint Selection Mechanism or the Object Selection Mechanism, is either a search waypoint or the location of a cargo, a fixed enemy emplacement, or home base. Search waypoints are generated from the search affordance map of the world, which indicates the degree to which each world location affords locomotion and search for objects via scout radar. The peaks in this affordance map constitute candidate search waypoints. The selection mechanism operates by finding an appropriate peak to serve as the current scout waypoint by combining the affordance related information with fuel maintenance and backtracking avoidance constraints.

TABLE 4

SCOUT ACTION MECHANISMS

<u>Mechanism Name</u>	:	Load/Unload Action Mechanism
Description	:	Activated by the Scout Action Mechanism to either load cargo or unload cargo and replenish resources at home base. Activates the Button Pressing Mechanism to perform the desired function.
Parameters	:	None

<u>Mechanism Name</u>	:	Enemy Action Mechanism
Description	:	When activated by the Scout Action Mechanism to attack an enemy craft, activates the Attack Monitoring Mechanism to guide the attack.
Parameters	:	None

<u>Mechanism Name</u>	:	Motion Action Mechanism
Description	:	For manual control crew models, activates the Manual Control Mechanism to move scout in a series of small movements to a desired waypoint (either a search waypoint, a cargo, a fixed enemy craft, or home base). For autopilot control model, activates the Autopilot Management Mechanism to command to autopilot to move the scout to the waypoint.
Parameters	:	Manual vs. Autopilot Scout Control
Values	:	Crew 1 Model : Autopilot Control
	:	Crew E Model : Manual Control
	:	Crew 2 Model : Manual Control

TABLE 4
(continued)

SCOUT ACTION MECHANISMS

Mechanism Name:	Attack Monitoring Mechanism
Description :	Activates Button Pressing Mechanism to fire missiles at an enemy until either the scout or the enemy is destroyed
Parameters :	None

Mechanism Name:	Autopilot Management Mechanism
Description :	For autopilot control model, activates the Cursor Mechanism to move the world display cursor to the specified waypoint (either a cargo, fixed enemy craft, a search waypoint, or home base), and then activates the Button Pressing Mechanism to initiate the autopilot.
Parameters :	None

Mechanism Name:	Manual Control Mechanism
Description :	For manual control crew models, accepts a a specified waypoint (either a cargo, a fixed enemy craft, a search waypoint, or home base) and generates a series of sub-waypoints to the waypoint. (See following pages for a description of this mechanism. Activates the Joystick Control Mechanism to simulate the physical joystick activities to fly to each sub-waypoint.
Parameters :	DL = Value of destination affordance as a percent of local affordances which determine path direction. (See text)
Values :	Crew 1 Model - Not used
:	Crew E Model - DL = 100%
:	Crew 2 Model - DL = 100%

The Scout Manual Control Mechanism also uses the search affordance world map to generate a path to the desired peak. Unlike the planning mechanism, though, which operates upon an global, high-level representation of this map in terms of peaks and ridges, the Scout Manual Control Mechanism uses the map in full detail in order to generate a path to the desired waypoint that is sensitive to local search affordance gradations. This mechanism operates by treating each world location as if it exerts an attractive force on the scout that is proportional to its search affordance, and inversely proportional to the cube of its distance from the scout. This inverse-cubed scaling factor is used to ensure that nearby locations have a larger attractive force than distant locations. In addition, the desired waypoint is considered to exert an attractive force. To guarantee that the waypoint will be reached, this force is independent of the distance between the waypoint and the scout.

Paths are generated by using vector addition to find the net attractive force acting upon the scout. The next step in the path is in the direction of this net force. The resulting direction of motion is thus determined by a combination of affordances from both the scout destination and the local environmental structure. In such a model, a much larger weighting on local affordances than on the affordance provided by the destination would result in somewhat aimless meandering though the world that is primarily determined by the local environmental affordance structure. If, on the other hand, a much larger weighting is used for the destination affordance than for the local affordances, nearly linear travel to the destination with insensitivity to the local environment would result.

In the model for describing Crews 2 and E, the value of the attractive force exerted by the destination was set equal to the net attractive force exerted by the local environment. Thus, for example, if the local affordance structure exerts a northward force on the scout, and the destination is to the east of the scout, the resulting direction of motion would be to the northeast. This equal weighting scheme has the possibility of stalling the scout due to force vectors that exactly cancel, but this phenomenon was never observed.

The final set of action mechanisms are those responsible for simulating the physical activity that is required to control the scout and friendly craft. All of the action mechanisms described above only organize interface activity, while the Physical Action Mechanisms appearing in Table 5 on the following page are actually responsible for simulating the execution of control actions. These action mechanisms are all quite simple, and the performance parameters for each crew model are summarized in the table. These parameters are approximate values measured from videotapes of the behavior of each crew.

TABLE 5

PHYSICAL ACTION MECHANISMS

Mechanism Name	:	Button Pressing Mechanism
Description	:	Simulates the physical activity of pressing a button on the control panel.
Parameters	:	TB = Button pressing time
Values	:	Crew 1 Model - TB = 1.50 seconds
	:	Crew E Model - TB = 0.75 seconds
	:	Crew 2 Model - TB = 0.75 seconds

Mechanism Name	:	Cursor Moving Mechanism
Description	:	Simulates the physical activity of moving the world display map cursor to a specified location. Uses Fitts' Law to estimate time required to move map cursor a distance D from previous position to a target of width W. D and W measured in display pixels (display width = 512 pixels = 30 cm). Width of cargo and enemy targets = 5 pixels, width of search waypoint target = 10 pixels. Movement time = A + B x Log(2D/W)
Values	:	Crew 1 Model - A = 2 ; B = .33
	:	Crew E Model - A = 1 ; B = .15
	:	Crew 2 Model - A = 1 ; B = .15

Mechanism Name	:	Editing Mechanism
Description	:	Simulates the physical activity of typing a friendly craft action command on the text editor.
Parameters	:	TE = Time to enter a mobile enemy attack command TO = Time to enter all other commands
Values	:	Crew 1 Model - TE = 4.0 ; TO = 9.0 Crew E Model - TE = 3.0 ; TO = 4.0 Crew 2 Model - TE = 3.0 ; TO = 5.0

TABLE 5
(continued)

PHYSICAL ACTION MECHANISMS

<u>Mechanism Name</u>:		Joystick Control Mechanism
Description	:	Simulates the physical activity of moving the scout manual control joystick to fly the scout along a path. Operates by moving the scout in a sequence of small steps. Step size is determined by the speed with which the scout is flown.
Parameters	:	SS = Scout speed (mph)
Values	:	Crew 1 Model - Mechanism not used
	:	Crew E Model - SS = 80.0
	:	Crew 2 Model - SS = 65.0

As mentioned previously, the timing parameters used in the physical action mechanisms represent the total time associated with both the preparation and execution of physical actions. For example, the mean button pressing time for Crew 1 was observed to be approximately 1.50 seconds, twice that of the mean times for Crews 2 and E. The reason for this large difference was not due to differences in the actual motion times associated with pressing buttons. Rather, this large timing difference results from the observation that Crew 1 took a longer time to identify appropriate buttons than did the other crews.

A similar difference between Crew 1 and the other two crews concerned the time required for preparing cursor manipulation actions. The larger A (or intercept) parameter value in the Fitts Law calculation used for describing Crew 1's behavior indicates that he took a longer fixed time to make cursor movements in addition to that time corresponding to the difficulty of the movement task. The larger B (or slope) parameter used for Crew 1 indicates that he was also more sensitive than the other two crews to the difficulty of the movement task. Approximate A and B parameters were derived by using a simple linear regression to fit two sessions of cursor movement time data. To represent the fact that movements to search targets were less constrained than movements to displayed objects, the width of search targets was set equal to twice the width of object targets.

Perceptual Mechanism Parameters

The perceptual mechanisms, shown in Figure 9 in the previous chapter, were parameterized by hypothesizing affordance structures that seemed to be consistent with the behavior of each of the three crews. The perceptual mechanisms dedicated to the tasks of search waypoint selection for the scout and friendly craft are described in Table 6 on the following pages. The first two mechanisms, the Scout Search Affordance and Scout Locomotion Affordance Mechanisms, were used in the models of Crews 2 and E to describe the manual control search activity of these two crews. The final four mechanisms, the Friendly Search Affordance, Friendly Collision Affordance, Friendly-Scout Coordination, and Friendly Range Affordance Mechanisms, were used in the models of all three crews to describe the selection of friendly craft search waypoints.

It was hypothesized that Crew 1's selection of search waypoints for the scout could also be described with the mechanisms for friendly craft waypoint selection. Two factors suggest this hypothesis. First, Crew 1's scout control behavior did not appear sensitive to the superior ability of the scout to discover objects with its large radar range. Crew 1 did not appear to differentiate between the scout and friendly craft for the purposes of loading cargo, thereby keeping the scout from its most important task of searching. Second, Crew 1's use of the autopilot

TABLE 6

SEARCH PERCEPTUAL MECHANISMS

Mechanism Name: Scout Search Affordance Mechanism

Description : Generates a search affordance world map used by Scout Waypoint Selection Mechanism to generate search waypoints and by Scout Manual Control Mechanism to generate paths between waypoints.

Parameters : SO = Search affordance value of open regions
SL = Search affordance value of lightly forested regions
SH = Search affordance value of heavily forested regions
SB = Search affordance value of regions beyond world boundaries
BT = Search affordance value of previously searched regions

The search affordance of a world point is calculated as follows. Let,

FO = fraction of world points within 1.5 mile radar range of the given point that are in open regions
FL = fraction of world points within 1.5 mile radar range of the given point that are in light forest
FH = fraction of world points within 1.5 mile radar range of the given point that are in heavy forest
FB = fraction of world points within 1.5 mile radar range of the given point that are beyond world boundaries

Then for unsearched points,

Search affordance = FOxSO + FLxSL + FHxSH + FBxSB

For previously searched points, Search affordance = BT

Values :

Parameter	Crew 1	Crew E	Crew 2
SO	Not used	0	0
SL	"	1	1
SH	"	1	1
SB	"	0	0
BT	"	-1.3	-1.3

(See text for a discussion of the parameter values)

TABLE 6
(continued)

SEARCH PERCEPTUAL MECHANISMS

<u>Mechanism Name</u>:		Scout Locomotion Affordance Mechanism
Description	:	Generates a locomotion affordance world map used by Scout Waypoint Selection Mechanism to generate search waypoints and by Scout Manual Control Mechanism to generate paths between waypoints.
Parameters	:	LO = locomotion affordance value of open regions LL = locomotion affordance value of lightly forested regions LH = locomotion affordance value of heavily forested regions LB = locomotion affordance value of regions beyond world boundaries
Values	:	

Parameter	Crew 1	Crew 2	Crew E
LO	Not used	0.0	0.0
LL	"	-1.5	-1.5
LH	"	-2.0	-2.0
LB	"	-3.0	-3.0

<u>Mechanism Name</u>:		Friendly Search Affordance Mechanism
Description	:	Generates a search affordance world map used by the Friendly Waypoint Selection Mechanism to generate search waypoints for the friendly craft in all three crew models, and for the scout in the Crew 1 model.
Parameters	:	FSO = search affordance value for open regions FSL = search affordance value for lightly forested regions FSH = search affordance value for heavily forested regions FSB = search affordance value for regions beyond world boundaries

TABLE 6
(continued)

SEARCH PERCEPTUAL MECHANISMS

Mechanism Name: Friendly Search Affordance Mechanism (continued)

Values :

Parameter	Crew 1	Crew E	Crew 2	Crew 1S
FSO	2	2	2	2
FSL	2	2	2	0
FSH	2	2	2	0
FSB	0	0	0	0

Mechanism Name: Friendly Collision Affordance Mechanism

Description : Generates a collision affordance world map used by the Friendly Waypoint Selection Mechanism to generate search waypoints for the friendly craft in all three crew models, and for the scout in the Crew 1 model.

Parameters : COLL = Collision affordance of a world point resulting in a collision
FREE = Collision affordance of a world point not resulting in collision

Values :

Parameter	Crew 1	Crew E	Crew 2	Crew 1S
COLL	0	0	0	0
FREE	1	1	1	1

Mechanism Name: Friendly-Scout Coordination Affordance Mechanism

Description : Generates a coordination affordance world map used by the Friendly Waypoint Selection Mechanism to generate search waypoints for the friendly craft in all three crew models

Parameters : COOR = Coordination affordance of a world point simultaneously coverable by scout radar and the radar of the friendly craft whose waypoint is being generated.
NCOR = Coordination affordance of all other world points

TABLE 6
(continued)

SEARCH PERCEPTUAL MECHANISMS

Mechanism Name: Friendly-Scout Coordination Affordance Mechanism (continued)

Values :

Parameter	Crew 1	Crew E	Crew 2	Crew 1S
COOR	0	3	3	Not used
NCOR	0	0	0	Not used

Mechanism Name: Friendly Range Affordance Mechanism

Description : Generates a fuel-range affordance world map used by the Friendly Waypoint Selection Mechanism to generate search waypoints for the friendly craft in all three crew models and to generate scout search waypoints in the Crew 1 model

Parameters : FOK = Range affordance of a world point to which the friendly craft can travel and still have enough fuel to return home, if the craft does not have enough fuel to last until the end of the mission.
FNO = Range affordance of all other world points

Values :

Parameter	Crew 1	Crew E	Crew 2	Crew 1S
FOK	1	1	1	1
FNO	0	0	0	0

C-3

for scout control required this crew to enter scout waypoints with the world display map cursor, in the same way that search waypoints are entered for the friendly craft. This requirement may have encouraged this crew to develop similar policies for scout and friendly craft waypoint selection. In general terms, the hypothesis is that Crew 1's use of a supervisory mode of control over all five craft resulted in a lessened sensitivity to the primary functional difference between the scout and friendly craft in terms of search capability. It should be made clear, though, that this causally stated hypothesis is not being evaluated here. What is specifically being examined is the hypothesis that Crew 1's processes for selection of scout and friendly craft activities were more similar than the processes used by Crews 2 and E. The parameters used for scout waypoint selection for Crew 1 appear under the heading Crew 1S.

The first two perceptual mechanisms described in Table 6 indicate that the same search affordance and locomotion affordance parameters were used to describe scout search activities for both Crews 2 and E. The parameters listed in the table are the same values that were used to produce the affordance maps appearing in Figures 11-14 in the previous chapter. Recall that the final affordance map that is used by the Scout Waypoint Selection and Scout Manual Control Mechanisms is created by superimposing (adding the values of) the locomotion and search affordance maps.

The parameters used in the Scout Search Affordance Mechanism and Scout Locomotion Affordance Mechanism for Crews 2 and E were determined by an extensive (but not exhaustive) search of the parameter space. The behavior of the Scout Manual Control Mechanism, which uses locomotion and search affordances to generate scout motion, was varied by changing these parameter values. The parameter values given in Table 6 provided the best fit to crew behavior. In addition, the resulting values seem to be in reasonable alignment with task goals and constraints.

It is critical to keep measurement scaling issues in mind when interpreting these parameter values. Due to the way in which these values are additively combined and then used to generate scout motion, the choice of both zero point and measurement unit for the affordance value scales are arbitrary. The zero and unit are arbitrary because the motion generating mechanism is only sensitive to the relative affordance values, or the gradations, in the total affordance map. A change in zero would simply raise or lower the affordance value of every world point by the same amount, resulting in no overall change in the relative affordance structure. A change of unit would alter the severity of the affordance gradations, but would once again leave the relative affordance structure unaltered. In the terminology used to describe the Scout Manual Control Mechanism, a change of unit would change the norm (or length) of the resultant force vector acting upon the scout but it would not change the direction of this vector. Since the direction of scout motion is determined by the direction of the resultant force vector, but left unaffected by the norm of this vector, a change in the affordance measurement unit does not affect the behavior of the Scout Manual Control Mechanism. Therefore, the only meaningful interpretations of the search and locomotion affordance mechanism parameters concern their relative values.

Purely as a matter of convention, then, search affordances were assigned positive values and locomotion affordances were assigned negative values. This arrangement was chosen to be consistent with a description of searching as behavior directed toward the (positively valued) goal of finding objects constrained by the (negatively valued) difficulty of locomoting through forested regions. The affordances for both searching and locomoting in open regions were assigned a value of zero, because open regions seemed to provide a neutral reference value to which the value of searching and locomoting in the forested regions could be compared. Assigning zero affordance values to open regions is not meant, for course, to indicate that locomoting through open regions was of zero difficulty or that searching open regions was of zero benefit in discovering objects. Rather, this zero point for the affordance scales was chosen purely as a matter of convenience.

The search affordances of both lightly and heavily forest regions were assigned equal and positive values to generate search behavior consistent with the forest boundary hugging behavior of Crews 2 and E (see Chapter V). The search affordance for previously searched points was given a negative value to lessen the amount of backtracking produced when traveling between search waypoints. This negative value also helped "smooth out" the scout paths by maintaining a large negative force directly behind the scout. The locomotion values for forested regions that were found to provide a good fit to crew behavior were indicative of the relative densities of the lightly and heavily forested regions. A large negative value for locomoting beyond the world boundary was used to keep scout paths within the displayed world.

The crews were hypothesized to be quite similar in terms of the affordances used to select friendly craft search waypoints, as indicated by the parameter values for the final four perceptual mechanisms in Table 6. These four mechanisms produce the affordance maps that are combined and used by the Friendly Waypoint Selection Mechanism. Whereas the final affordance map for the scout is produced by simple addition of affordances, the final affordance map for the friendly craft is produced by addition and multiplication. Note that the Friendly Collision and Friendly Range Mechanisms both produce zero-one binary valued maps. The affordance values of these two maps are each multiplied with the sum of the other two maps, the Friendly Search Affordance and Friendly-Scout Coordination Affordance maps. This process is used to exclude a search waypoint from consideration if its assignment to a craft would result in a collision or a craft running out of fuel.

The calculations used to generate the affordance maps indicating collision points and the locus of points within fuel range produce exact and correct values for these measures. It is almost surely the case that crews did not make exact and veridical estimates of fuel ranges and collision points, even though no crew ever had two craft collide or a craft run out of fuel in the sessions used for analysis. Since exact calculations were not performed, good performance was probably the result of processing mechanisms that were conservatively biased. That is, friendlies were probably not sent to the edges of their fuel ranges, and any pair of friendlies were probably kept separated by more than the distance necessary to avoid collision. Unfortunately, no mechanisms could be developed to operate in a fashion consistent with these behaviors, other than by simply adding a conservative tolerance to the equations used in the exact calculations. This increase in complexity would not increase the psychological plausibility of these mechanisms, though, since the mechanisms would still based on exact (albeit conservative) calculations. Therefore, it was decided to use the simpler, exact calculations in order to produce behavior in high-level agreement with crew behavior. The issue of identifying psychologically plausible mechanisms to perform fuel range estimation and collision avoidance was not addressed further.

Unlike the scout searching affordances, the parameter values in the Friendly Search Affordance Mechanism are not related to forest density. This is due to the fact that the speed of friendly craft motion is not affected by forest density when the craft fly above tree level. Although crews differed with respect to their ability to keep the friendly craft above tree level, the friendly craft were kept above the trees for the majority of time by each of the three crews. Therefore, the affordance values were based on the assumption that the friendly craft would fly above tree level and thus be unaffected by forest density.

The only difference between the three crews in terms of friendly search perceptual mechanism parameters concerns the friendly-scout coordination affordances. Since Crew 1 was hypothesized to not differentiate between the scout and friendly craft, the model for this crew was not designed to give any extra weighting to friendly craft search waypoints that would enable coordination with the scout.

The parameters for describing scout waypoint selection in the Crew 1 model differ from the friendly waypoint selection parameters in two respects. First, since Crew 1 confined scout activities to open regions (presumably due to slow autopilot travel through forests), the search affordance value of forested regions was set equal to zero. In addition, since the scout cannot coordinate its activities with itself, the coordination affordance was not used to determine scout waypoint affordances.

The final set of perceptual mechanisms to be discussed concern sensitivity to affordances pertaining to loading cargo, attacking enemy craft, traveling to home base, avoiding collisions, and maintaining craft altitude. Since the parameter values in these mechanisms were hypothesized to vary greatly between the crews, the method that was used to select the parameter values should be described. As discussed previously in this chapter, the modeling approach is based on the assumption that the role of the perceptual mechanisms is to provide information that is used to implement, but not determine, a particular strategy. This assumption is based on the framework developed for describing highly skilled performance. In the suggested framework, strategic issues are properly discussed at the teleological level where task goals are considered to determine how a particular strategy could be developed off-line, or incrementally, by a process of learning or reasoning. It was argued that, for modeling purposes, these issues can be kept distinct from the issues concerning how a given strategy is implemented in real-time.

Since the current modeling addresses only the implementational issues, it may seem as if the process of parameterizing these perceptual mechanisms is shallow or descriptive. That is, the parameterization approach is based on considering what different sets of affordance values might be sufficient to generate the high-level, strategic differences that were observed in the experiments. The selected parameter values are thus responsible for implementing a given strategy, but they do not indicate how a particular strategy might have evolved through learning.

The final set of perceptual mechanisms are described in Table 7 on the following pages. Different parameter values in the cargo affordance mechanism were used to adjust for the degree to which the three crews differentially used the scout to load cargo. The parameter value used for Crew 1 results in identical cargo affordance calculations for both the scout and friendly craft. The parameter values for the other two crews result in lower cargo affordance values for the scout relative to the friendly craft. Given a scout and friendly craft with identical distances to a given cargo, fuel levels, numbers of missiles, and weight carrying capacities, the scout would have a lower affordance for loading the cargo than would the friendly craft.

In both the Cargo Affordance Mechanism, the distance between the craft and the cargo (DC) enters the affordance calculation in the following way (see Table 7). The factor

TABLE 7

OBJECT PERCEPTUAL MECHANISMS

Mechanism Name: Cargo Affordance Mechanism

Description : Generates the cargo-loading affordance used by the Object Selection Mechanism to select scout and friendly craft activities.

Parameters : SC = The ratio of scout to friendly craft cargo affordances. A low value of SC results in a low affordance for scout cargo loading relative to the friendly craft affordance. A unit value for SC results in no difference between the scout and friendly craft affordance calculations.

The affordance for a cargo-craft pair is calculated as follows:

Let,

DC = the distance between the cargo and the craft
UN = 1 if cargo can be loaded and unloaded at home base without the craft running out of fuel or the mission terminating,
= 0 otherwise
MS = 1 if the craft has at least 2 missiles,
= 0 otherwise
WC = 1 if the craft has enough weight carrying capacity remaining to load the cargo,
= 0 otherwise
FE = 1 if the cargo is not within 0.4 miles of a fixed enemy craft,
= 0 otherwise

Then, if the craft is a friendly, the affordance is:

$$\mathrm{MAX}\ [\ 0,\ ((1 - DC/7) \times UN \times MS \times WC \times FE)\ ]$$

And, if the craft is the scout, the affordance is:

$$\mathrm{MAX}\ [\ 0,\ ((1 - DC/7) \times UN \times MS \times WC \times FE \times SC)\ ]$$

Thus, the affordance is greatest when the DC is small, and is equal to zero when DC is at least 7 miles.

Values :

Parameter	Crew 1	Crew E	Crew 2
SC	1	0.33	0.67

TABLE 7
(continued)

OBJECT PERCEPTUAL MECHANISMS

Mechanism Name: Fixed Enemy Affordance Mechanism

Description : Generates the fixed-enemy affordance used by the Object Selection Mechanism to select scout and friendly craft activities.

Parameters : FF = A number in the interval [0,1] that determines friendly craft affordance from the "nominal" fixed enemy affordance.
SF = A number in the interval [0,1] that determines scout affordance from the "nominal" fixed enemy affordance.

The nominal affordance for a enemy-craft pair is calculated as follows:

Let,

DF = the distance between the enemy and the craft
UN = 1 if enemy can be attacked without the craft running out of fuel or the mission terminating,
= 0 otherwise
MS = 1 if the craft has at least 3 missiles,
= 0 otherwise

Then, the nominal affordance (NOM) is:

NOM = MAX [0, ((.85 - .85xDE/7) x UN x MS)]

Thus, the nominal affordance is greatest when DE is small, and is equal to zero when DE is at least 7 miles.

For the friendly craft, the affordance is: NOM x FF

For the scout, the affordance is: NOM x SF

Values :

Parameter	Crew 1	Crew 2	Crew E
FF	0	1	1
SF	1	0.33	0.33

TABLE 7
(continued)

OBJECT AFFORDANCE MECHANISMS

<u>Mechanism Name</u>:		Home Affordance Mechanism
Description	:	Generates the home affordance used by the Object Selection Mechanism to select scout and friendly craft activities
Parameters	:	CH = Gain on nominal home affordance due to need to unload cargo

The home affordance is the maximum of three nominal home affordances: the affordances for unloading cargo (NOMC), for refueling (NOMF), and for resupplying missiles (NOMM).

The three nominal affordances are calculated as follows:

Let,

XSF = The percent of current fuel that is in excess of the amount of fuel needed to safely return home, given that the current fuel level is not sufficient to last until the end of the mission.

WCR = The weight carrying capacity remaining in pounds [0 < WCR < 1000].

MIS = the number of missiles carried [0 < MIS < 6]

Then, the nominal home affordance for refueling is:

NOMF = 1.0 - XSF/100

The nominal home affordance for cargo unloading is:

NOMC = (1000 - WCR/1000)/1000

And the nominal home affordance for resupplying missiles is:

NOMM = 1 if MIS is 0 or 1; 0 otherwise

The home affordance is:

MAX [1, MAX(NOMF,NOMM,(CH x NOMC))]

Values :

Parameter	Crew 1	Crew 2	Crew E
CH	3	1	1

TABLE 7
(continued)

OBJECT AFFORDANCE MECHANISMS

<u>Mechanism Name</u>: Mobile Enemy Affordance Mechanism

Description : Generates affordance for attacking mobile enemies used by Object Selection Mechanism to select scout and friendly craft activities.

Parameters : None

The affordance value for the friendly craft equals 1.01 when the craft is locked-on by mobile enemy craft radar. The affordance value for the scout equals 1.01 when the scout is below tree level and locked-on by mobile enemy craft radar. These are the situations in which the craft cannot typically escape from the attacking enemy craft.

<u>Mechanism Name</u>: Collision Affordance Mechanism

Description : Generates affordance for taking a collision avoidance action for a friendly craft. Used by the Object Selection Mechanism to select scout and friendly craft activities.

Parameters : None

The affordance value equals 1.00 when a craft is expected to collide with another friendly craft. Taking an avoidance action with the first of the two relevant craft zeros the collision affordance for both craft.

<u>Mechanism Name</u>: Friendly Up Affordance

Description : Generates affordance for assigning an up action for a friendly craft. Used by the Object Selection Mechanism to select scout and friendly craft activities.

Parameters : None

The affordance is zero for all craft above tree level and for craft below tree level in open regions. A value of 0.1 is incrementally added to the affordance level each time the map display is viewed for craft down in forests, up to a maximum of 1.0

(1 - DC/7) is used to make the affordance for loading a piece of cargo a function of the distance between the craft and the cargo locations. The distance is divided by seven to normalize distances between seven and zero miles to a zero-one range. Although the maximum possible distance between a craft and a cargo is approximately 14 miles, the zero value on the scale is set to seven miles to increase the sensitivity of the affordance calculation to small differences in distances. The maximization operation in the affordance calculation ensures that a negative affordance value is never produced. A similar distance calculation is used in the Fixed Enemy Affordance Mechanism.

The Fixed Enemy Affordance Mechanism was similarly tuned to the varying degrees to which the three crews seemed to differentiate between the scout and friendly craft. Given a nominal affordance level as described in the table, the affordance values for each crew for both the scout and friendly craft were expressed as functions of this nominal value. The nominal affordance was defined to the (same) affordance used for the scout in the Crew 1 model and for the friendly craft in the models of Crews 2 and E. Since Crew 1 did not attack fixed enemy craft with friendly craft, this nominal affordance was nulled in this model for these craft. In addition, since Crews 2 and E did not often use the scout to attack fixed enemy craft, this affordance value was lessened for the scout in these two models. The maximum possible fixed enemy affordance value produced by this mechanism is 0.85, whereas the maximum possible cargo affordance value is 1.00. This difference was required due to the fact that less points were scored by destroying a fixed enemy than were scored by unloading a piece of cargo at home base. Therefore, loading cargo was preferred to attacking fixed enemy craft.

The home affordance parameter values were also tailored to individual crews. The home affordance due to the need to unload cargo used in the Crew 1 model was greater than in the other two crew models to produce behavior consistent with Crew 1's policy of serially processing cargo with the scout and friendly craft. As discussed previously in this chapter, the maximization operation in the Home Affordance Mechanism was used to generate a thresholding of the affordance values. The final three perceptual mechanisms, the Mobile Enemy, Collision, and Up Affordance Mechanisms were identical in the three crew models.

The affordances produced by these perceptual mechanisms provide the input to the selection mechanisms described in the previous chapter. As mentioned above, the same selection mechanisms were used to model all three crews. Thus, the parameter values given above for both the perceptual and action mechanisms were, with one exception, the only adjustments made to fit the behavior of the human crews. The final parameter that was manipulated was the time required for the perceptual mechanisms to operate. Up to this point, the only mechanisms that have been assumed to require processing time have been the action mechanisms that simulate observable control activities at the interface. It is obvious that, in the terminology of the present model, the operation of the crew's perceptual and selection mechanisms cannot be instantaneous. Unfortunately, there are no directly measurable data that could be used to estimate this time.

Therefore, the processing time due to the operation of the perceptual and selection mechanisms was estimated by examining the behavior of the model with different timing parameters. The need to consider this time parameter surfaced during the Crew 1 model fitting process. For this crew especially, it was found that the model produced far too little idle time for the friendly craft. The amount of idle time produced by the Crew 1 model that was due solely action mechanism delays was roughly one half of the craft idle time observed for Crew 1. Therefore, some perceptual or central delays had to be assumed in order to match the behavior of this crew. Perceptual or central delays also had to be assumed for the Crew 2 and Crew E models as well, although these were much smaller than the delay used for mimicking Crew 1. The perceptual mechanism processing time parameters that seemed to provide best fits to the data were: Crew 1 model - 3 seconds; Crew 2 model - 1 second;

Crew E model - 0.5 seconds. These times reflect the amount of time required by the model to update the array of affordances each time the model "looked" at the world display. These parameter values were selected because they provided the best fits to crew performance.

Table 8 on the following pages provides a summary of the parameter values that were manipulated in the attempt to mimic the performance differences between the three crews. Only those parameters that were varied to fit the crews are included in the table. This concludes the discussion of the model parameterization process.

TABLE 8

SUMMARY OF PARAMETER VALUES

1. Mechanism: Scout Action Mechanism
   Parameter: Serial vs. Parallel scout & friendly control

   Crew 1: Serial
   Crew E: Serial
   Crew 2: Parallel

2. Mechanism: Fixed Enemy Action Mechanism
   Parameter: Fixed enemy action command

   Crew 1: Not used
   Crew E: E:^,G,A,A,A
   Crew 2: E:^,G,A,A,A

3. Mechanism: Search Action Mechanism
   Parameter: Search action command

   Crew 1: S:^,G
   Crew E: S:^,G,P
   Crew 2: S:^,G

4. Mechanism: Home Action Mechanism
   Parameter: Home action command

   Crew 1: X:U (no cursor required)
   Crew E: X:U (no cursor required)
   Crew 2: S:^,G,v,U (cursor required)

5. Mechanism: Motion Action Mechanism
   Parameter: Manual versus Autopilot scout control

   Crew 1: Autopilot Control
   Crew E: Manual Control
   Crew 2: Manual Control

6. Mechanism: Manual Control Mechanism
   Parameter: Destination affordance as a percent of local affordances.

   Crew 1: Not used
   Crew E: 100%
   Crew 2: 100%

TABLE 8
(continued)

SUMMARY OF PARAMETER VALUES

7. Mechanism: Button Pressing Mechanism
   Parameter: Button Pressing time

   Crew 1: 1.50 seconds
   Crew E: 0.75 seconds
   Crew 2: 0.75 seconds

8. Mechanism: Cursor Moving Mechanism
   Parameter: Fitts Law intercept

   Crew 1: 2 seconds
   Crew E: 1 second
   Crew 2: 1 second

   Parameter: Fitts Law slope

   Crew 1: 0.33 seconds/bit
   Crew E: 0.15 seconds/bit
   Crew 2: 0.15 seconds/bit

9. Mechanism: Editing Mechanism
   Parameter: Time to enter a mobile attack command

   Crew 1: 4.0 seconds
   Crew E: 3.0 seconds
   Crew 2: 3.0 seconds

   Parameter: Time to enter all other commands

   Crew 1: 9.0 seconds
   Crew E: 4.0 seconds
   Crew 2: 5.0 seconds

10. Mechanism: Joystick Control Mechanism
    Parameter: Maximum scout speed

    Crew 1: Not used
    Crew E: 80 m.p.h
    Crew 2: 65 m.p.h

TABLE 8
(continued)

SUMMARY OF PARAMETER VALUES

Mechanism: Friendly Search Affordance Mechanism
Parameter: Search affordance of light forests

Crew 1:	2
Crew E:	2
Crew 2:	2
Crew 1S:	0

Parameter: Search affordance of heavy forests

Crew 1:	2
Crew E:	2
Crew 2:	2
Crew 1S:	0

Mechanism: Friendly-Scout Coordination Affordance Mechanism
Parameter: Affordance of regions simultaneously coverable by scout and friendly craft

Crew 1:	0
Crew E:	3
Crew 2:	3
Crew 1S:	Not used

Mechanism: Cargo Affordance Mechanism
Parameter: Ratio of scout to friendly craft affordance

Crew 1:	1.00
Crew E:	0.33
Crew 2:	0.67

Mechanism: Fixed Enemy Affordance Mechanism
Parameter: Ratio of friendly craft affordance to nominal affordance

Crew 1:	0
Crew E:	1
Crew 2:	1

TABLE 8
(continued)

SUMMARY OF PARAMETER VALUES

Mechanism: Fixed Enemy Affordance Mechanism
Parameter: Ration of scout affordance to nominal affordance

Crew 1:	1.00
Crew E:	0.33
Crew 2:	0.33

Mechanism: Home Affordance Mechanism
Parameter: Gain on home affordance due to need to unload cargo at home base

Crew 1:	3
Crew E:	1
Crew 2:	1

Mechanism: All Perceptual Mechanisms
Parameter: Updating time

Crew 1:	3.0 seconds
Crew E:	0.5 seconds
Crew 2:	1.0 seconds

Model Evaluation

The performance measures described in Chapter III that were used to construct crew profiles were also used to construct profiles from the behavior produced by the model with each of the three parameter sets. The model was run on each of the same eight world-configuration/payoff-structure combinations used in the three crew experiments. Thus, the model was run on each world configuration twice, once with the payoff favoring cargo processing and once with the payoff favoring enemy processing. Since the model is not sensitive to payoff structure information, the model should, theoretically, produce the same stream of behavior and world events in the two runs using the same world configuration.

Here, unfortunately, implementation details serve to cloud the issue. Since the computer system on which the model ran was not dedicated to this modeling task, there was some chance that the execution of other users' tasks could slightly alter the time required for one of the mechanisms to process information. For example, in one run of the model a certain cargo was discovered a fraction of a second before a different cargo was discovered, but in the second model run, this sequence was reversed. This would, in turn, alter the friendly craft assignments to these cargo, completely changing the locus of the paths traced out by the two friendly craft in traveling to these cargo. This change would then result in differences in the times at which the remaining cargo were discovered and enemy craft were encountered. Within a few minutes after the minor divergence between the two model runs, the two models would be operating on highly different worlds and producing entirely different streams of behavior.

To reduce the severity of such effects, the model was run at hours when there were no or very few other users. Even, though, when the model and simulation were the only tasks running on the computer system, there was typically some divergence in model behavior between successive runs, probably due to external or historical factors that altered the way in which the many computer programs comprising the model were selected for processing by the cpu scheduler. This scheduler is responsible for determining the order in which the many programs awaiting cpu time were executed. Certain factors appeared to contribute to these scheduling operations that were not controllable by the computer user. For example, the scheduler would sometimes perform system maintenance operations during the time in which the model was running, thereby slightly altering the way in which the model's programs were selected for execution. For analysis purposes, then, the actual stream of behavior produced by the model in any given run is perhaps viewed as a somewhat "accidental" property of the model, whereas the relevant outputs of the model are the behavioral invariants that can be discovered in higher level performance assessments.

The same performance measure calculations were used to analyze the data produced by both the model and the human crews. The degree to which the model and crew profiles agreed was used as a measure of the empirical adequacy of the model for each crew. As in the crew profile comparisons, a t-test was used on each performance measure to provide a measure of similarity between model and crew behavior. Due to the fact that the model could produce very similar behavior in each pair of sessions using the same world configuration, a correlational analysis of the data produced by the model was performed to determine the appropriate degrees of freedom that should be used to analyze model performance. The only measures on which the model's performance were found to be highly correlated between the two sessions on the same world configuration were the measures of area searched by the scout and the total area searched by the models of Crews E and 2. For the statistical tests using these measures, then, the degrees of freedom were reduced to account for the fact that the intra-world variability produced by the model was very low.

Two types of similarity tests were performed. The first type measured to similarity between each human crew and the model used to mimic that crew. For each of the 22 performance measures in the profiles, t-tests were used to identify whether crew and model performance differed, using the same significance level as was used in the comparisons between the three human crews. These tests will be referred to as similarity tests in the following chapter where the modeling results are discussed.

This method of measuring model adequacy is problematic, of course, due to the fact that high variability in model performance can serve to render all comparisons between crew and model performance statistically insignificant. To guard against this possibility, each of the three crew models were also compared against each other, in the hope of obtaining the same pattern of differences between the three models as was observed between the three crews. These tests cannot be artificially passed due to high variability in model performance, and will be referred to as configural tests in the following chapter. It should be pointed out, though, that the variability in performance on almost every one of the 22 measures was lower for the model than for the human crews. This finding strengthens the validity of the similarity testing procedure, and suggests that running additional model sessions to improve the power of the statistical tests is of limited utility.

The use of certain assumptions that are known to be false presents some problems for the model evaluation approach. The first such assumption concerns the simplified set of action commands used in the model to control the friendly craft. As mentioned above, crews used a range of different action commands for essentially the same purpose, although the rationale behind command selection could not be hypothesized. The inability to determine policies for use of the "up" action is one important example. Crews 1 and 2 usually prefaced search, cargo, and fixed enemy attack actions with "up" actions, but this was not always the case. Since consistent policies for "up" action usage could not be found, the behavior of these crews was approximated in the model by using the action commands that these crews used most frequently: those containing "up" actions. Therefore, the performance of the model might be expected to differ with the performance of Crews 1 and 2 with respect to the measure of the percent of time spent above trees for the friendly craft. Note that it should be possible to achieve a better match with the data on this measure by including a pseudo-random component in the model that would selectively employ "up" actions with the same relative frequency as used by the crews, although this gain in empirical adequacy would not provide any additional insight into the reasons why crews differentially used the "up" actions in their friendly craft command strings.

Additional issues associated with the use of false assumptions will be discussed as they arise in the discussion of modeling results in the following chapter. The example discussed above has been presented in some detail to convey the complexity of the problems associated with attempting to validate a generative model, which is known to be simplified, by comparing its behavior with the behavior of human subjects.

A similar yet related problem is trying to isolate just those empirical data that should be used to evaluate the model, since the assumptions of the modeling approach are only intended to apply to a restricted range of empirical phenomena, namely, highly skilled performance. It became clear that, even with the most highly skilled crew, situations arised for which the crew would attempt to generate novel solutions to familiar problems, in an attempt to improve on steady-state strategies. Unfortunately, the behavior associated with these "experiments" is embedded within the stream of behavior produced by the mechanisms that were hypothesized to produced highly skilled perceptually-based behavior. Unless these data can be teased out, there is a danger of penalizing the model for failing to deal with phenomena for which it was not designed to account.

On the other hand, there is also a possibility of inappropriately crediting the model for producing empirically adequate behavior when that behavior is generated "for the wrong reasons." For example, the model may produce behaviorally valid results due to a fortuitous cancellation of the effects of erroneous assumptions. It is natural during the model fitting process for attention to be drawn to mismatches between model and human behavior, often allowing for model deficiencies that do not directly produce inappropriate activity to go unnoticed.

As a final note concerning the validity of comparing model and crew behavior, one adjustment made to the dynamics of the controlled system must be discussed. As mentioned above, the model used for mimicking the behavior of Crew 2 did not attempt to scatter the waypoint destinations of friendly craft around home base. It was hypothesized that Crew 2 scattered these waypoints to lessen the probability of two craft colliding while returning home. (None of the three crews had two craft collide in the sessions used for analysis). The model contained two provisions for avoiding collisions. The first was to exclude search waypoint assignments that would result in expected collisions. The second was the use of a perceptual affordance mechanism specifically designed to identify potential collisions. The affordance value produced by this mechanism would signal the action mechanisms to implement an evasive action with one of the friendly craft.

Even though the model contained these provisions for collision avoidance, it was found that craft collisions occurred in the sessions performed by the model with too high a frequency. The major reason for this problem seemed to concern the limited ability of the model to employ self-referential reasoning. When the model was evaluating which potential friendly craft waypoints could result in collision, the model had to rely upon knowledge of the time at which the friendly craft would be expected to begin its travel to the waypoint. But this time depended upon the time at which the model "itself" would be able to determine and implement a course of action. This time, of course, was partially determined by whether the waypoint that was currently under consideration would be free of collision, or whether additional waypoints and other actions had to be explored, evaluated, and compared. And whether or not the current waypoint under consideration would cause a collision was determined by . . ., and so on. This difficulty was compounded by the fact that the mechanism computing the waypoint affordance did not have access to other contextual information that impacted the time at which the friendly craft would initiate travel to the waypoint. Since no original solution to this problem was forthcoming in this research, the solution that was adopted was to radically decrease the distance between two craft that would cause a collision. Therefore, the model did not attempt to scatter waypoints around home base in an attempt to avoid a collisions between friendly craft. This alteration in the collision distance calculation was the only change made to the dynamics of the system to accommodate the model.

CHAPTER VII

MODELING RESULTS

Introduction

In this chapter, the results of crew modeling will be discussed. As described in the previous chapter, two types of tests were performed to assess the empirical adequacy of the model for each crew. Similarity tests were performed to identify how well the parameterized models produced behavior in agreement with the crew performance profiles described in Chapter III. Statistical tests were used to identify any differences between crews and their associated models on each performance measure. Since success on these tests can be inappropriately achieved due to high variability in model performance, configural tests were also performed to assess how well the performance differences between each pair of crews were also exhibited by each associated pair of crew models. These configural tests are more stringent measures of the empirical adequacy of the crew models. At the end of the chapter, general conclusions pertaining to inter-crew differences and model adequacy will be discussed.

<u>Summary of Results</u>

This section is a general summary of the results of both the similarity and configural tests. After this overview, the adequacy of the crew models with respect to each of the 22 performance measures in the crew and model profiles will be discussed in more detail.

A total of 66 similarity tests were performed (3 crew-model pairs each using 22 measures). In 58 of the tests, model and crew performance could not be differentiated. All crew models satisfied tests associated with points scored, discovering cargo and enemy craft, and successfully processing cargo and enemy craft. Also satisfied were all tests associated with the differential usage of the scout and friendly craft pertaining to these measures. On the other hand, the model failed to achieve consistency with crew performance on eight of the 66 measures. (Only 3.3 failures would be expected due to chance if the measures were independent: 66 tests x .05 alpha level. More failures would be expected if the measures were not independent). Of these failures, one occurred for the model of Crew 1, three occurred for the model of Crew 2, and the model of Crew E failed on four performance measures. All three models failed on the measure of the time spent for friendly craft above tree level; in each of these cases the difference between model and crew performance was approximately 13%. As will be discussed below, two of the Crew E model failures pertaining to searching performance appear to be attributable non-stationarity in Crew E behavior that was not captured by the model.

For analysis of the results of the more demanding configural tests, four types of test results were defined. A "hit" was defined as an agreement between the comparison tests for each pair of crews and the associated crew models on a specific performance measure. For example, if Crew 1 sighted less cargo than Crew 2, and the Crew 1 model sighted less cargo then the Crew 2 model, this test was scored a hit. Another type of hit is when neither the crews nor the associated crew models differed on a measure. A "miss" was defined as a difference exhibited between a pair of crews on a specific measure that was not exhibited by the associated pair of crew models on that measure. A "false alarm" was defined as a difference that was found between a pair of crew models that was not exhibited by the associated pair of crews. Finally, a "reversal" was defined as a case where crews differed in one direction, whereas the associated pair of crew models differed in the other direction. For example, if Crew 1 loaded more cargo than Crew 2, but the Crew 1 model loaded less cargo than the Crew 2 model, this result was termed a reversal.

As with the similarity tests, a total of 66 configural tests were performed (3 pairs x 22 performance measures). Of these 66 tests, 54 were hits, five were misses, six were false alarms, and one test resulted in a reversal. Three of the misses occurred in the comparisons of the models of Crews 1 and E, and the other two misses concerned the comparisons of the models of Crews 2 and E. With all five misses, the model pairs and crew pairs differed in the same direction, although the difference in model performance was not great enough to achieve significance.

Of the six false alarms, three occurred in the comparisons of the models of Crews 1 and 2, and the other three misses occurred in the comparisons of the models of Crews 1 and E. With all six false alarms, the model pairs and crew pairs differed in the same direction, although the difference in model performance was great enough to achieve significance whereas the difference in crew performance was not. The reversal concerned the comparison of the time spent above trees by the friendly craft for Crews 2 and E and their respective models.

Modeling Results

In this section, the results of the similarity and configural tests will be discussed in greater detail. Figure 17 on the following page indicates the average number of points scored per session and the average number of enemy craft destroyed per session for each of the three crews and each of the three crew models. At the top of each graph, any significant differences between the performance of each

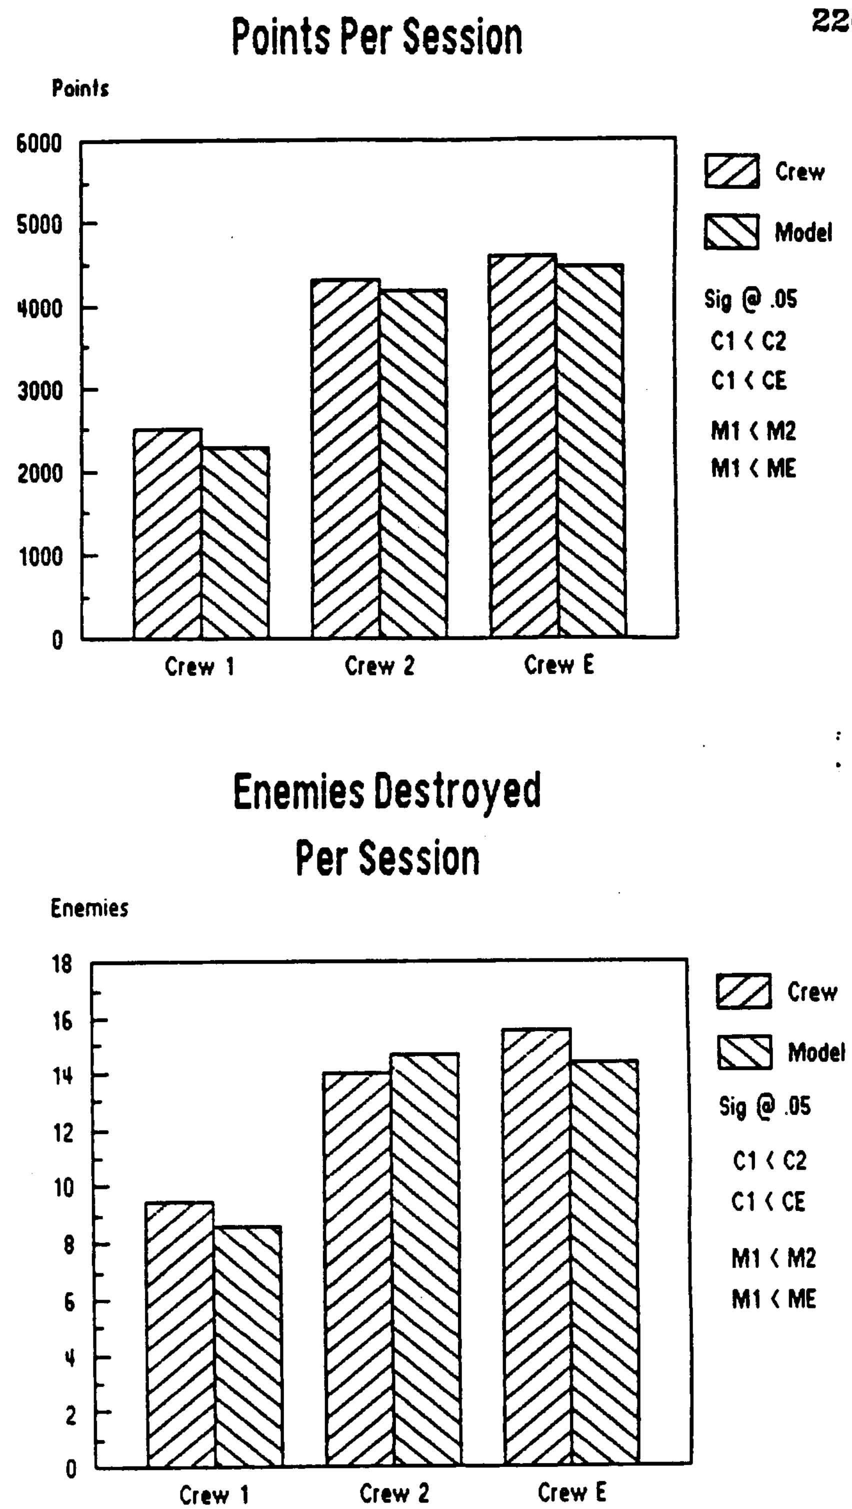

Crew and model comparisons of points and total number of enemy craft destroyed

Figure 17

pair of crews and each pair of crew models are noted. In later graphs where there are similarity test failures, any significant differences between each crew and its respective model will be noted as well. As can be seen in Figure 17, the three crew models satisfied all three similarity tests with respect to the measures of points scored and enemy craft destroyed. In addition, the configural differences between the three subjects all also replicated by the three models.

Figure 18 on the following page shows crew and model performance concerning the number of enemy craft destroyed by the scout and friendly craft. Once again, no model differed from its respective crew on either of these measures. On the other hand, Figure 18 represents the first configural test failure, a miss in this instance. Crew 1 destroyed significantly less enemy craft with the scout than did Crew E (3 versus 5.45 enemy craft). The model of Crew 1, on the other hand, did not destroy significantly less enemy craft with the scout than did the model of Crew E, although a similar trend was exhibited (3.25 versus 4.02 enemy craft).

Figure 19 on the next page indicates the total number of cargo unloaded at home base and the percent of discovered cargo that were eventually unloaded at home base. Once again, all similarity tests were satisfied with respect to these two performance measures. On the other hand, the

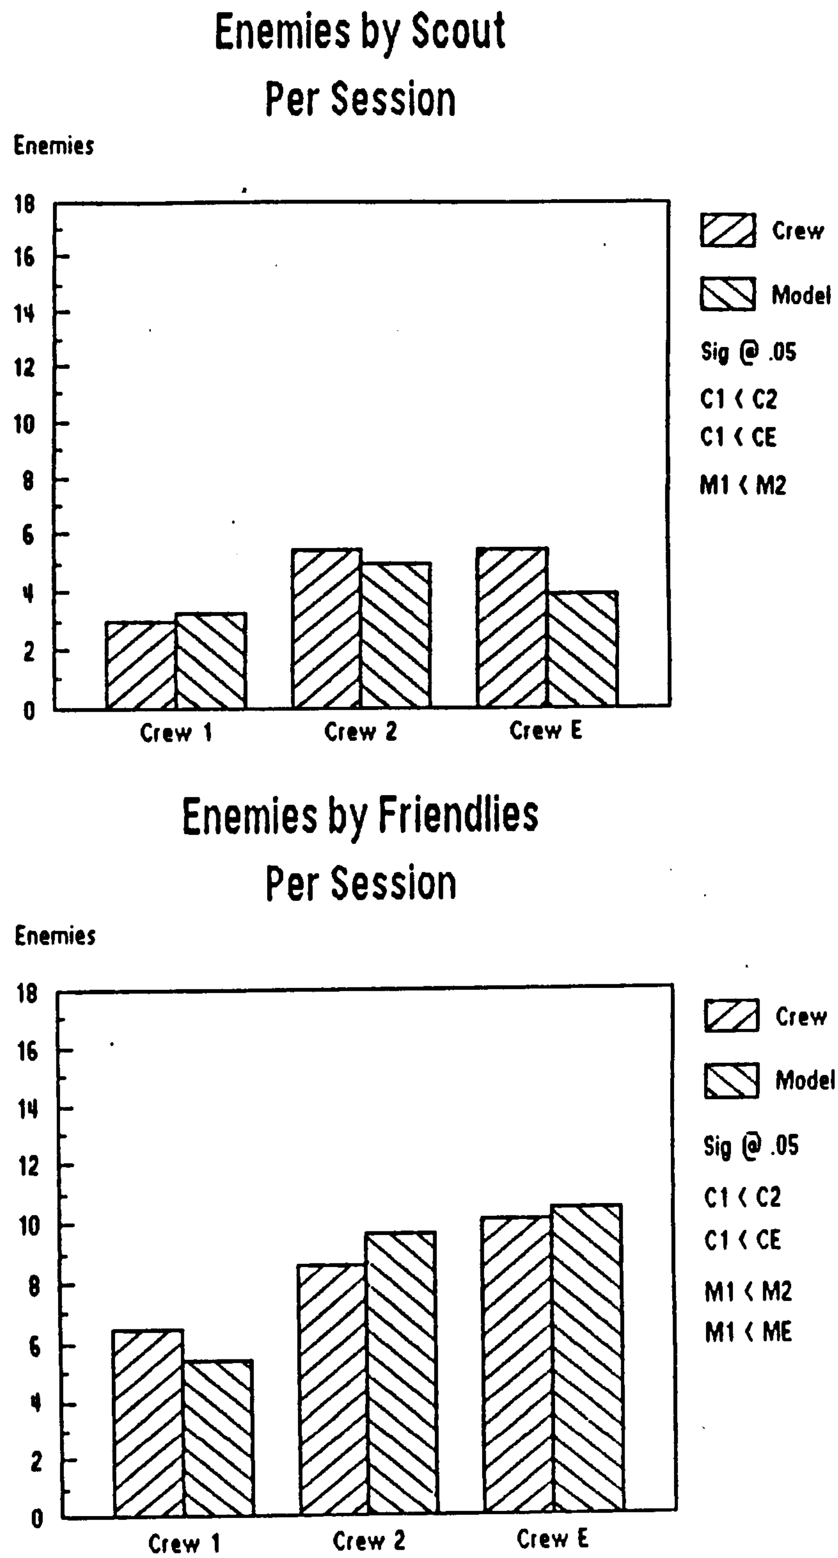

Crew and model comparisons of enemy destroyed by scout and friendly craft

Figure 18

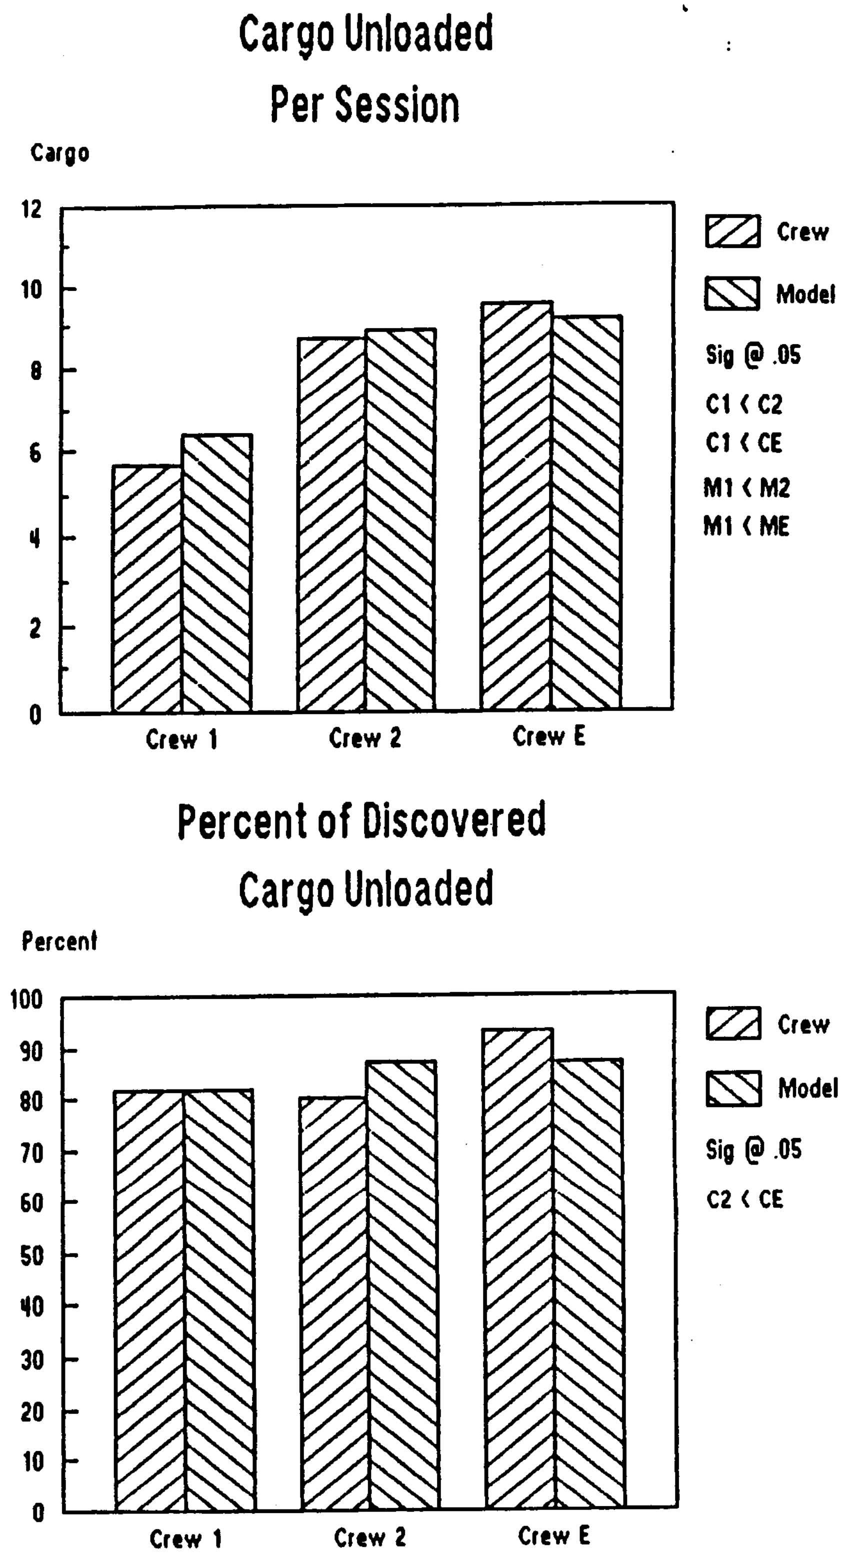

Crew and model comparisons of cargo unloaded and percent of discovered cargo unloaded

Figure 19

difference between Crews 2 and E on the measure of the percent of discovered cargo returned to home base was not replicated by the models of Crews 2 and E. The reason that the similarity tests were satisfied but one configural test was failed on this measure seems to be that the two crew models performed at a level half-way between the performance of the two crews. Thus, the two models were similar enough to their respective crews on this measure, but not sufficiently different to replicate the small but significant difference between these two crews.

The fact that the two models did not differ on this measure is due to the fact that no parameters were altered between these two models in an attempt to generate different performance on this measure. This measure represented one of the few differences between Crews 2 and E. The major phenomena that was selected to be explained concerning these two crews was how they were able to achieve such similar performance given their difference in crew size. The Crew E model was able to achieve the performance of the Crew 2 model despite the fact that its action mechanisms were constrained to perform consistently with assumptions concerning the peripheral limitations associated with the one-person crew condition. That this similarity in performance was achieved without replicating the superior ability of Crew E to unload a higher percentage of discovered cargo suggests that this superiority was not a major factor contributing to

Crew E's ability to perform at the level of a two-person crew.

Figure 20 on the following page indicates the average number of cargo unloaded by the scout and friendly craft for each of the three crews and each of the three crew models. All six similarity tests and all six configural tests were satisfied with respect to these two performance measures.

Figure 21 on the next page indicates the average number of friendly craft destroyed and the average number of cargo discovered for each of the three crews and three crew models. All six similarity tests pertaining to these two measures were satisfied. Only five of the six configural tests were satisfied, though, as the comparison between the models of Crew 1 and Crew E yielded a false alarm with respect to the average number of craft destroyed. No significant difference on this measure occurred between Crew 1 and Crew E, although Crew 1 averaged 1.29 craft destroyed per session whereas Crew E averaged 0.57 craft destroyed per session. The model of Crew 1 averaged 1.38 craft destroyed, and the model for Crew E averaged 0.38 craft destroyed per session, the difference between the two being statistically significant.

The failure of the model to achieve the appropriate configural results on this measure is particularly difficult to interpret. The author's extensive experience at this task suggests that the number of friendly craft destroyed

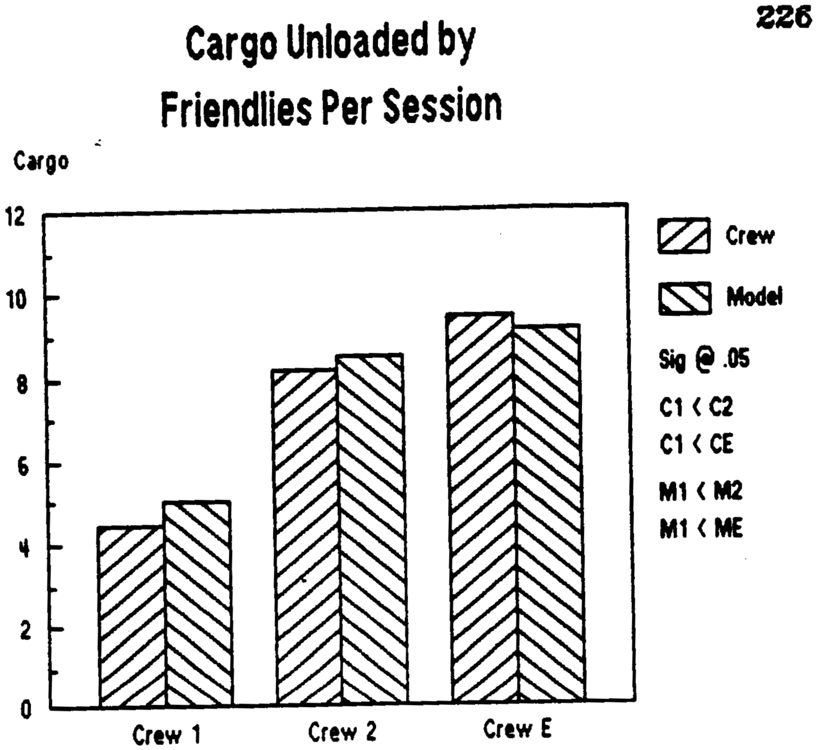

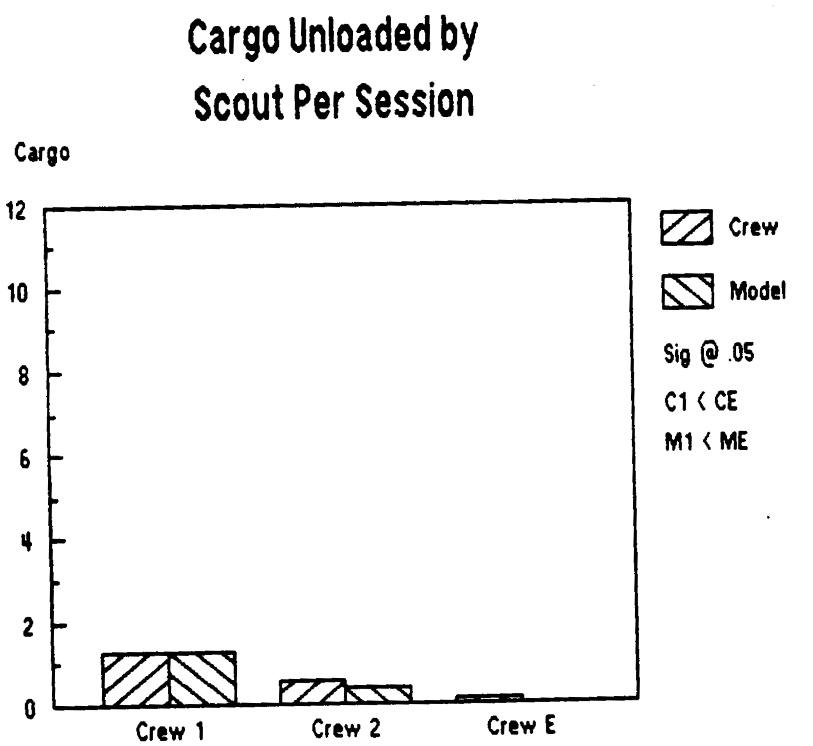

Figure 20

Crew and model comparisons of cargo unloaded by scout and friendly craft

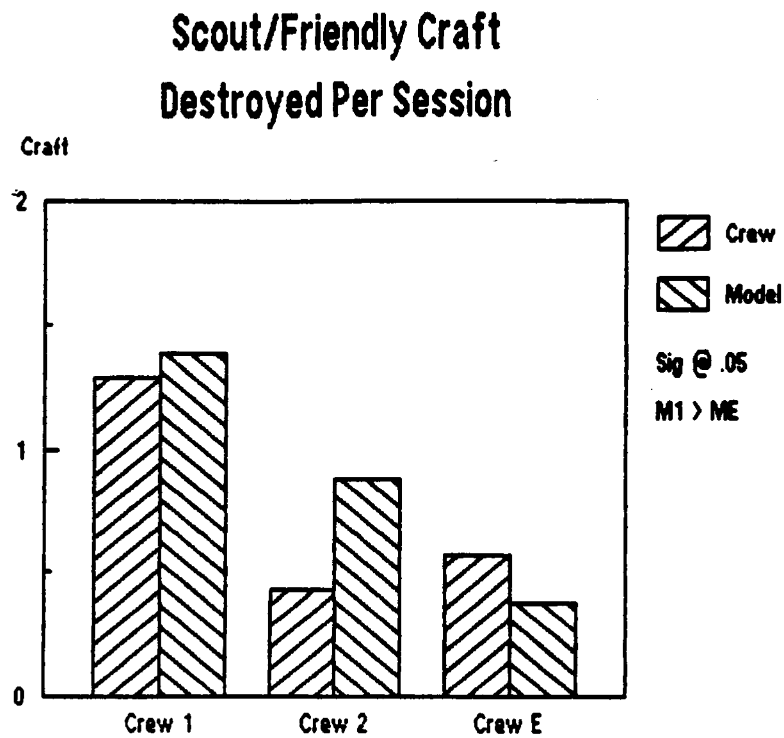

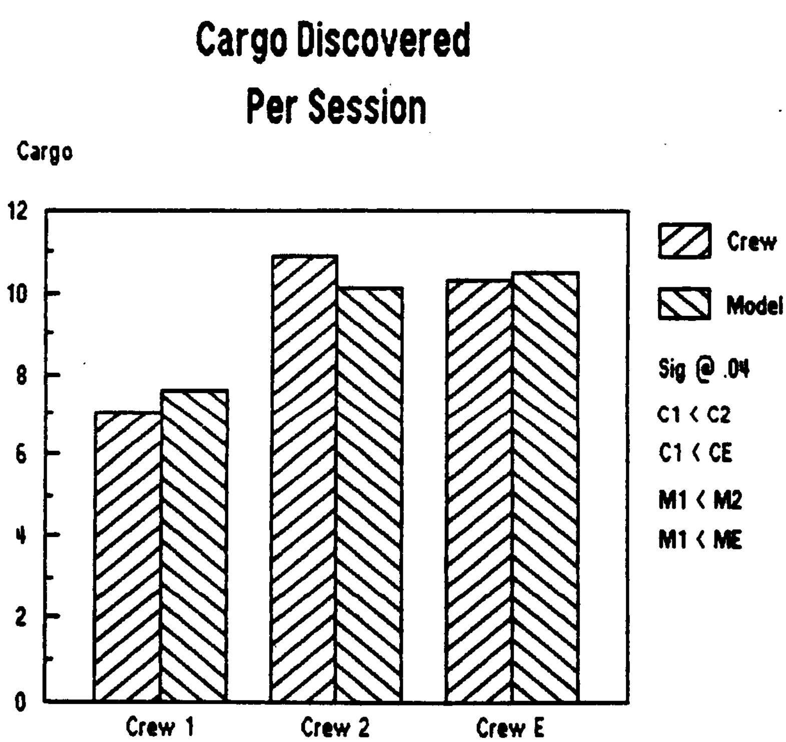

Figure 21

Crew and model comparisons of friendly craft destroyed and cargo discovered has a very large effect on the number of points scored. In addition to the 600 points that are deducted, the craft is not available to be used to earn points for the rest of the mission. The fact that Crew 1 averaged more than twice as many craft destroyed as Crew E, while being statistically insignificant, almost surely contributed to the statistically significant difference between the two crews in terms of total points scored. It would be almost impossible to average as many craft destroyed as did Crew 1 and still score as many points as did Crew E.

Figure 22 on the following page indicates the average number of cargo discovered by the friendly craft and the scout. All six similarity tests pertaining to these two performance measures were satisfied. With respect to the configural comparisons, one of the six tests resulted in a miss. Crew 1 sighted significantly less cargo with the scout than did Crew E (2.3 versus 5.1 cargo). The model for Crew 1 sighted less cargo with the scout than did the model for Crew 2 (3.2 versus 5.3), but this difference was not significant.

Figure 23 on the next page indicates the average time duration between loading and unloading cargo, and the average number of times the craft returned to home base to unload and refuel. All six similarity tests pertaining to these two measures were satisfied. On the other hand, four false alarms were evidenced in the configural tests due to

Cargo Discovered by Friendlies Per Session

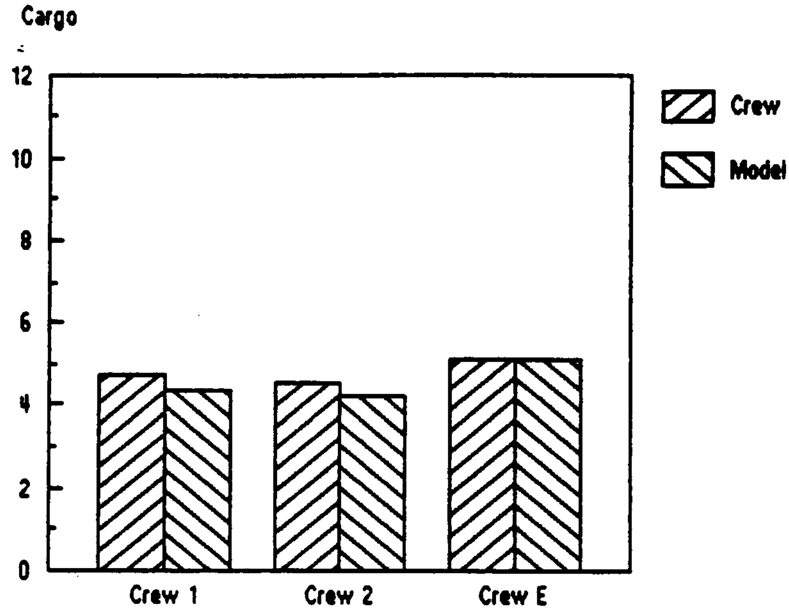

Cargo Discovered by Scout Per Session

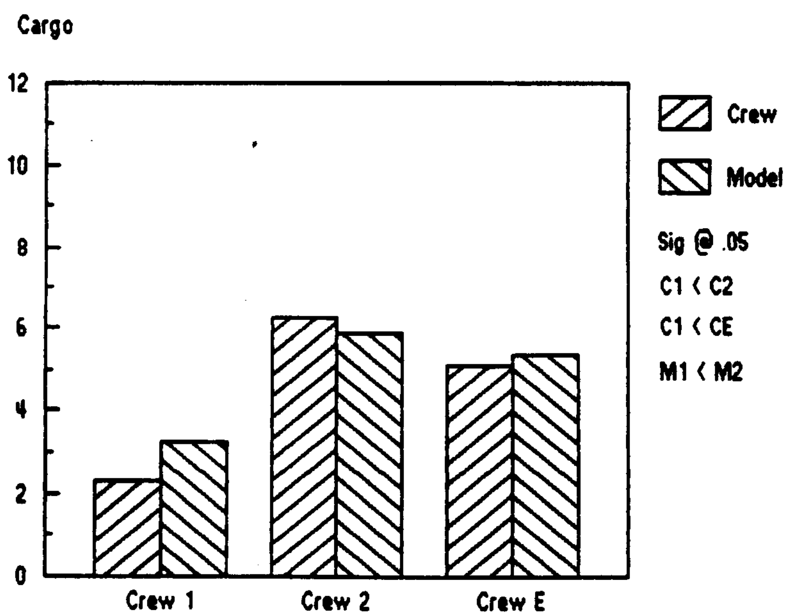

Figure 22

Crew and model comparison of cargo discovered by scout and friendly craft

Time Between Cargo Loaded and Unloaded

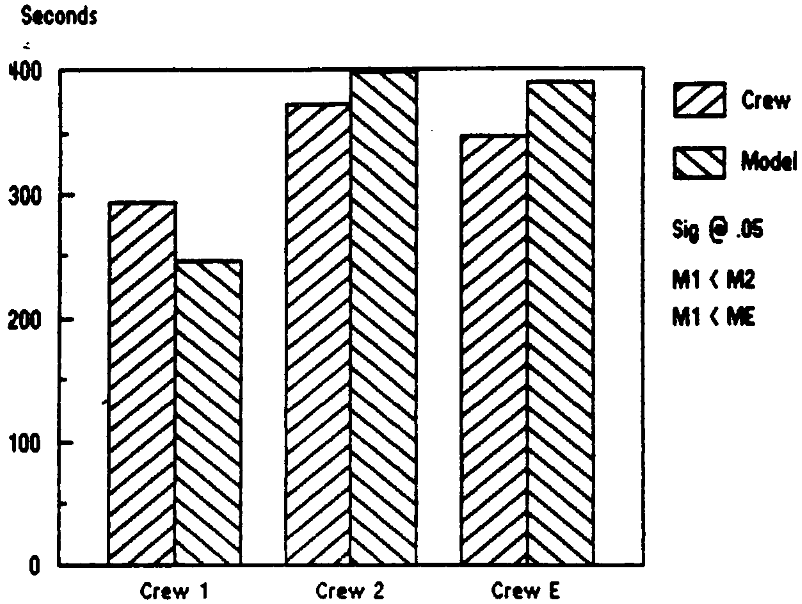

Number of Unloads at Home Per Session

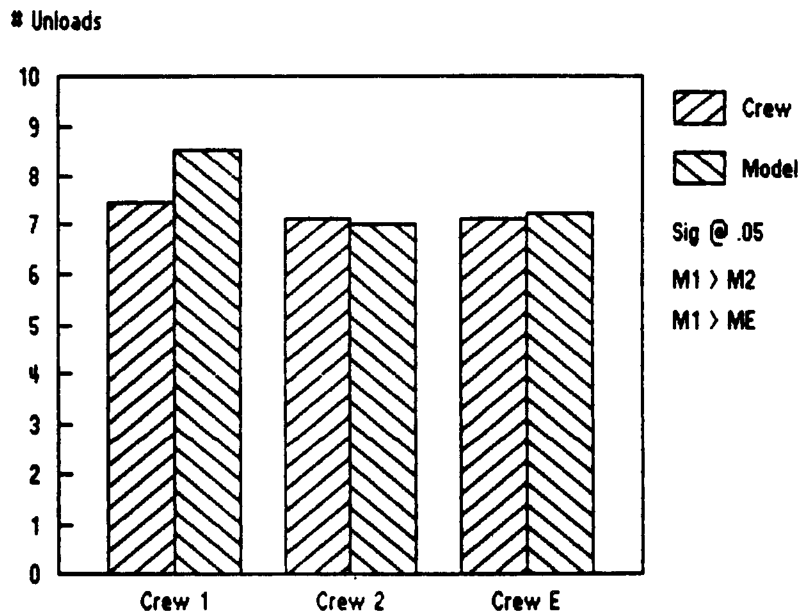

Figure 23

Crew and model comparisons of duration between loading and unloading cargo and the number of times craft sent home these measures. With respect to the duration between cargo loading and unloading, the difference between the Crew 1 model and Crew 2 model, and the difference between the Crew 1 model and the Crew 3 model were significant, whereas the corresponding differences in Crew performance were not. An examination of the session by session data indicated that the Crew 1 model had a much lower variability on this measure than Crew 1. Although the mean durations were similar enough so that the Crew 1 model did not differ significantly from Crew 1 on this measure, the much lower variance for the model resulted in significant differences between the Crew 1 model and the other two models.

The fact that the Crew 1 model had a shorter average duration between the times at which cargo were loaded and unloaded was caused primarily by the way in which this model serially processed cargo. Crew 1 was observed to have this sub-optimal policy, therefore the home affordance parameter reflecting the need to go home to unload cargo was raised in this model to produce this behavior. This serial cargo loading-unloading policy probably also resulted in the two false alarms pertaining to the number of times craft were sent home to unload. The Crew 1 model sent friendlies home significantly more often than did the models of Crew 2 and Crew E, whereas there was no such significant difference between the three crews.

What these four false alarms seem to suggest is that Crew 1 was less rigid in his use of the serial cargo loading-unloading policy than was the model of Crew 1. That is, the consistent employment of the serial policy in the model of Crew 1 was probably exaggerated with respect to the degree to which Crew 1 actually followed this policy. The Crew 1 model, constrained to perform rigidly in adherence to this policy, produced more highly stereotyped behavior regarding cargo unloading than did Crew 1, who may have been a bit more flexible in his adherence to this policy. This overly-stereotyped behavior from the Crew 1 model resulted in less variability with respect to the performance measures in Figure 23, thereby resulting in four false alarms in the configural tests.

This overly stereotyped behavior was caused in large part by the high weighting on the need to return loaded cargo to home base in the calculation of the home action affordance in the model of Crew 1. A large weight was required to match Crew 1's performance on the measure of the average (low) number of cargo unloaded per trip home by the friendly craft (see below). Lowering this weight resulted in cargo processing behavior that was less stereotyped and therefore resolved the mismatches on the measures mentioned above, but it caused the model of Crew 1 to unload too many cargo per trip home by the friendly craft. It therefore appears as if Crew 1's cargo processing behavior could not be adequately simulated by a mechanism that used a single parameter to determine when the friendly craft should be sent home to unload cargo.

Figure 24 on the following page indicates the average number of cargo unloaded per trip home by the scout and friendly craft. The six similarity tests and six configural tests pertaining to these performance measures were all satisfied.

Figure 25 on the next page indicates the average time spent idle for the scout and friendly craft. For the scout, the similarity tests for Crews 2 and E and the models of Crews 2 and E were not satisfied. These failures are due primarily to the difficulties that were encountered in trying to adequately interpret the craft idleness measure. For the crews, idleness was measured in the following way. Data were available concerning the position of each of the craft, including the scout, every 10 seconds. To calculate craft idleness, the number of 10-second intervals was counted in which a craft moved less than a threshold distance (representing 20 m.p.h). This number was divided by the total number of such intervals in the session (179 intervals). For the scout, the resulting number was used as an estimate of the percent of time that the scout was not being actively controlled. This number is a poor estimate due to the grain size of the discetization used and the fact that scout dynamics allowed the scout to glide an ap-

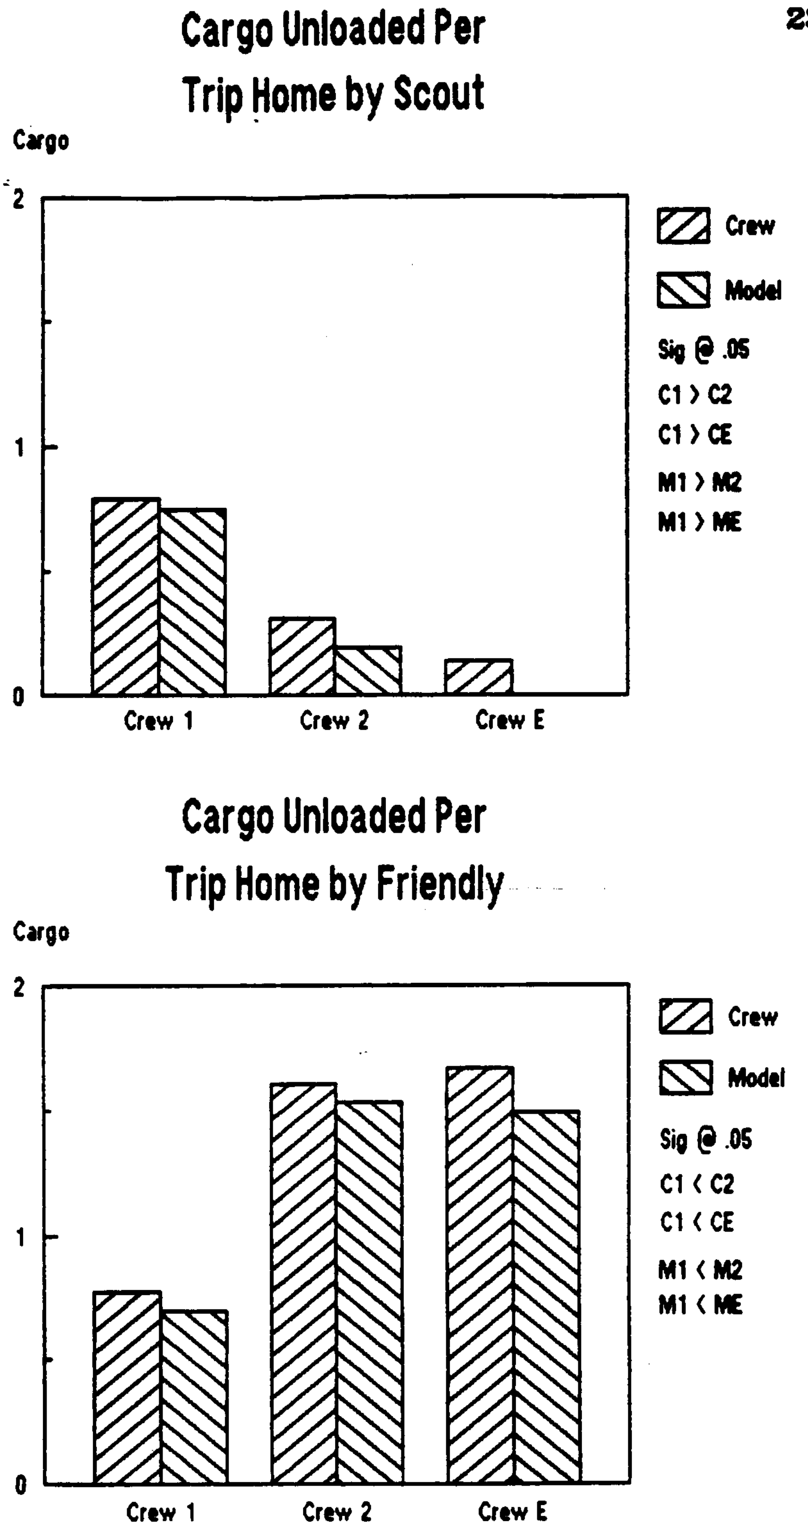

Figure 24

Crew and model comparisons of the number of cargo unloaded per trip home for the scout and friendly craft

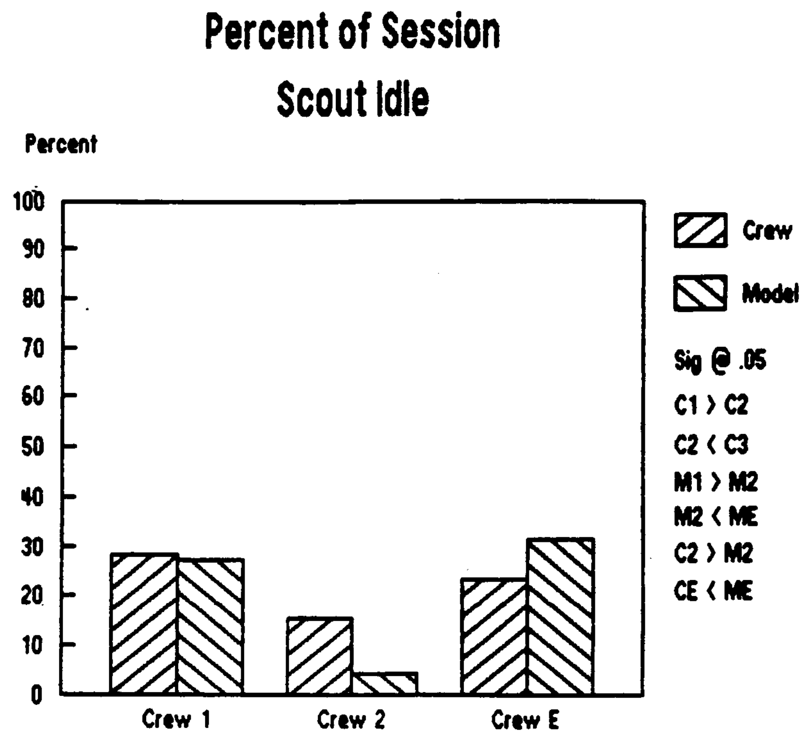

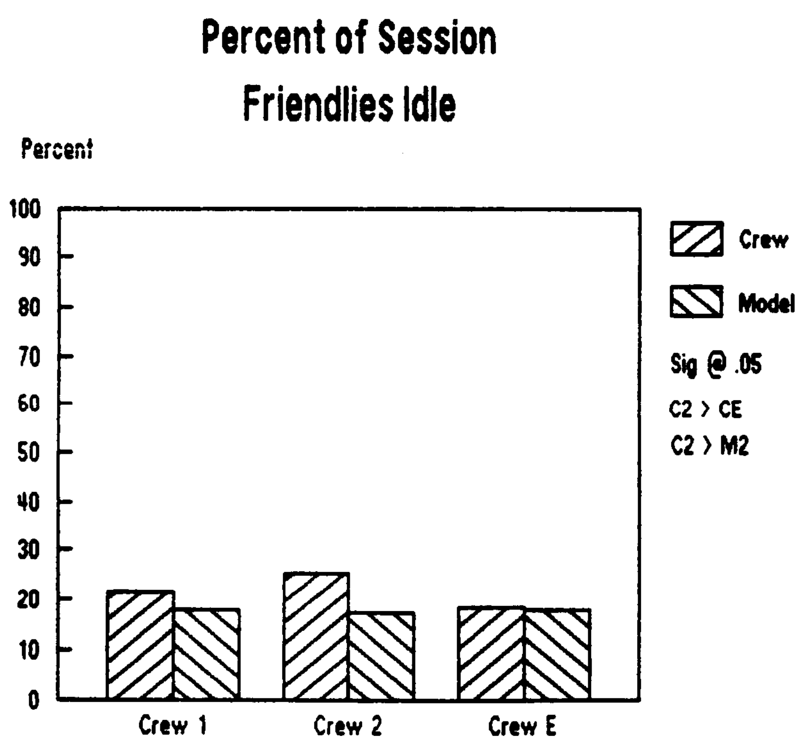

Figure 25

Crew and model comparisons of idleness

preciable distance even after active control via the joystick was stopped.

The assumptions employed in the scout manual control action mechanisms in the models of Crew 2 and E were simplified in the following way. The model of Crew 2, who had one crew member dedicated to scout manual control, had no provisions for less than active control of the scout. That is, no reasons were ever hypothesized to assume that the Crew 2 pilot was ever doing anything by controlling the scout. Thus, the scout, under control of the Crew 2 model, was never actually idle. In order to compare Crew 2 and model performance on this measure, then, the scout idleness measure had to be interpreted not as genuine lack of control effort, but rather as the inability of the Crew 2 pilot to maintain acceptable scout speed.

The fluctuations in scout speed due to encounters with world objects (e.g. trees, home base) were not produced by the model due to the use of a simplified mechanism for scout locomotion. Therefore, the scout control behavior produced by the model could not be tested against the idleness measure used to analyze crew performance. On the other hand, it was imperative that the amount of scout motion produced by the models of Crews 2 and E be consistent with the behavior of these crews. Therefore, the average distances traveled by the scout per session for Crews 2 and E were measured. Crews 2 and E did not differ on this measure. For the Crew 2 model, then, the average speed of travel for the scout was calculated under the assumption that the pilot was always performing scout motion control. This average speed was used as a parameter in the scout action mechanisms, as described in the previous chapter. Therefore, the average distance traveled by the scout while under control of the Crew 2 model was the same as when the scout was under control of Crew 2.

The Crew E model could not be parameterized in the same way since Crew E was not assumed to be able to perform continuous scout control due to the need to enter action commands for the friendly craft via the text editor. Therefore, videotapes were consulted that indicated that, when Crew E was performing scout motion control, the joystick was typically deflected to the maximum degree possible (80 m.p.h.). This speed was then used as a parameter in the scout action mechanisms. Note, though, that the scout in the Crew E model was only moved at this speed during the time intervals in which the model was not performing control activity related to the friendly craft. An analysis of the average distance traveled by the scout under control of the Crew E model yielded the same average distance as measured for Crew E.

Thus, the actual amount of scout control activity exhibited by the models of Crews 2 and E seemed to have been consistent with control activity exhibited by these crews.

The measures of scout idleness that appear in the previous figure have been calculated with the same measures used for the crews, although for reasons discussed above, these measures are not very meaningful. They have only been included to maintain consistency with the measures used in Chapter III for the analysis of crew performance.

The idleness measures for the friendly craft in Figure 25, on the other hand, are meaningful since friendly craft locomotion was produced in the same way for both the crews and the crew models. With respect to this measure, all tests were satisfied except for the fact that the Crew 2 model produced less idleness than did Crew 2 (18% versus 25%). When the Crew 2 model was adjusted to produce less friendly craft idleness by tuning the time required to update the perceptual mechanisms, the model performed worse than Crew 2 on measures dealing with altercations with enemy craft (too many friendly craft were lost).

This problem may reflect a serious deficiency of the model. Namely, all perceptual mechanisms were assumed to require the same time for updating, and this assumption may be unreasonable. A more realistic assumption would be that those mechanisms pertaining to lock-ons with enemy craft might be updated more quickly than, for example, the mechanisms responsible for detecting affordances associated with cargo. This difference might be due to the time-criticality of reacting to encounters between friendly and enemy craft, and the presence of an audio alarm that signaled lock-ons. This problem did not arise in the process of fitting the model to Crew 1, since the same (slow) perceptual updating time was consistent with both high measures of craft idleness and high numbers of friendly craft destroyed. The problem did not surface in modeling Crew E, since the same (fast) perceptual updating time was consistent with both low measures of craft idleness and low numbers of friendly craft destroyed by enemies.

Figures 26 and 27 on the following pages indicate the percentage of mission time spent above trees by the friendly craft and the effectiveness with which the simulated world was searched. The reason for the failures in the altitude tests were discussed in the previous chapter, where it was noted that the use of a simplified set of action commands for friendly craft control would require the average time spent above trees for the friendly craft to be incorrect. Fortunately, model performance did not seem to diverge from crew performance on this measure to a great extent.

All tests were satisfied with respect to the amount of world searched in total (Figure 26). For the amount of world searched with the scout and friendly craft (Figure 27), the test failures concern the model for Crew E. Crew E searched more of the world with the scout than did the Crew E model (65% versus 54%). Crew E also searched less of the world exclusively with the friendly craft than did the Crew

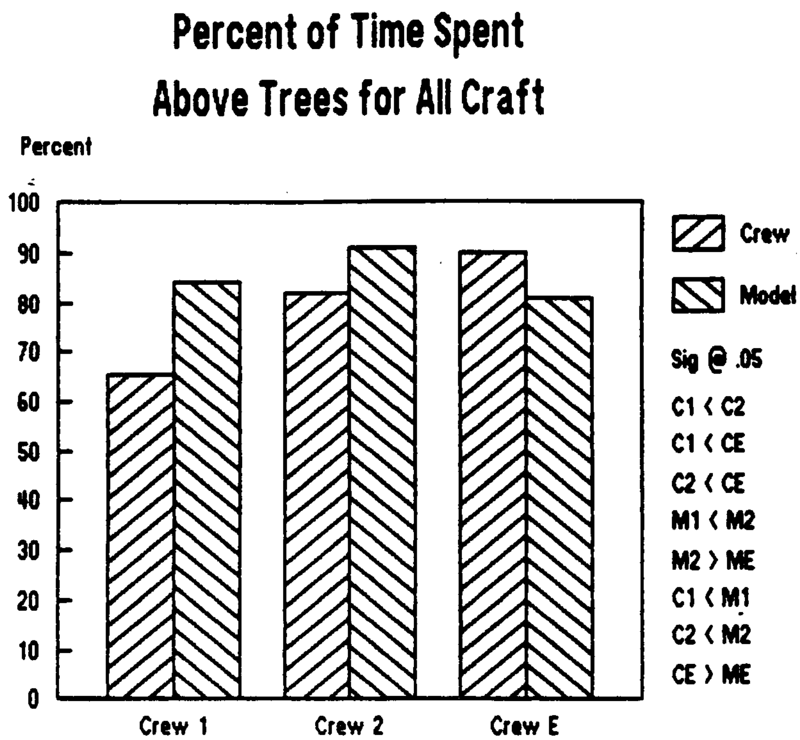

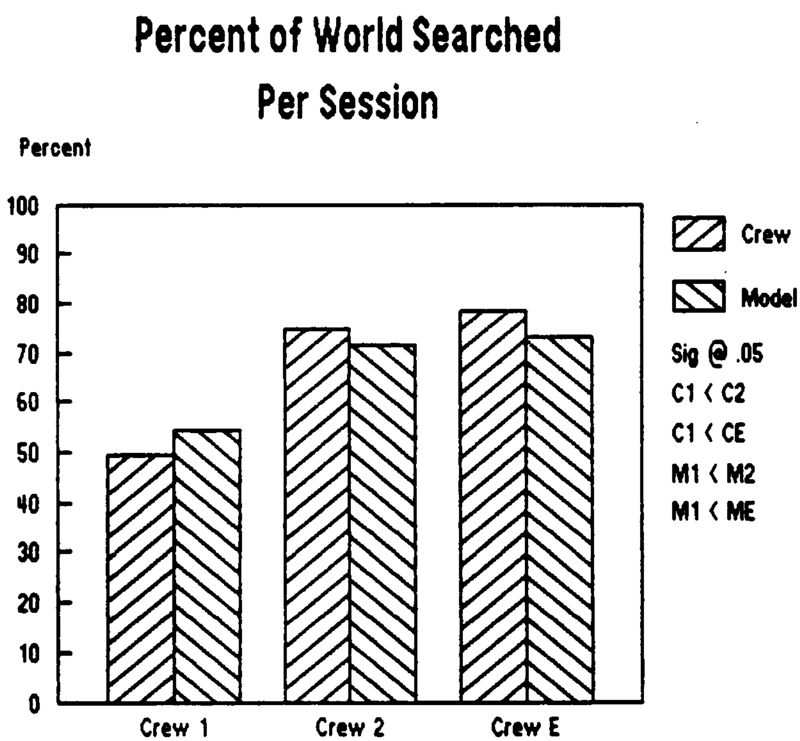

Figure 26

Crew and model comparisons of time spent above trees and total world searched

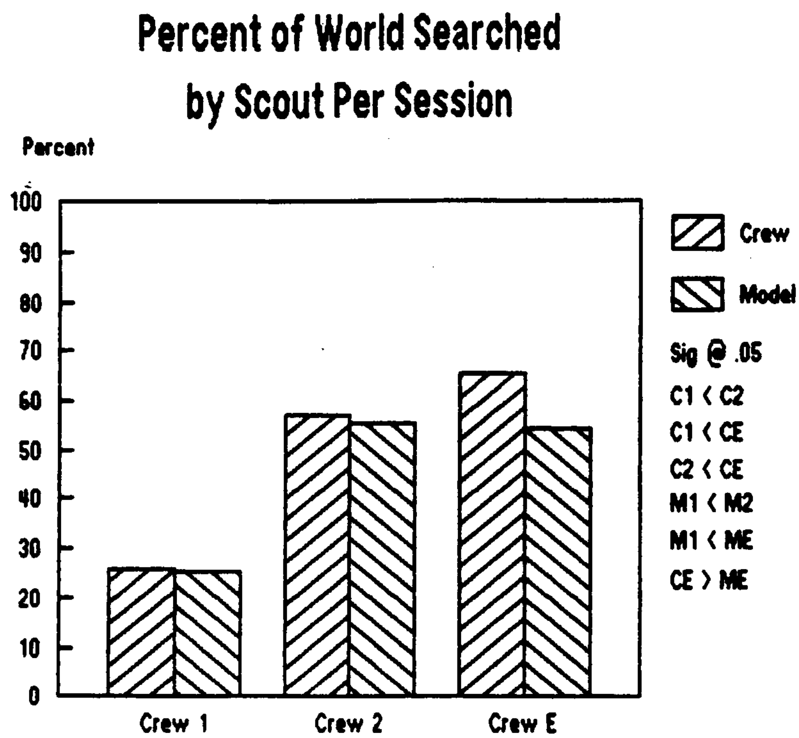

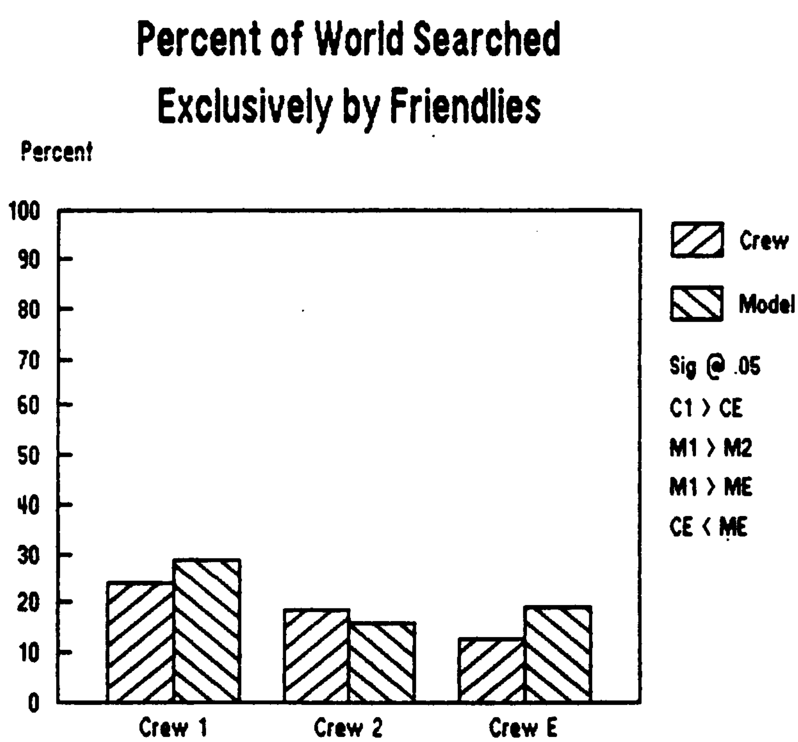

Figure 27

Crew and model comparisons of world
searched by scout and friendly craft

E model (12% versus 19%). Finally, the configural test between the models of Crews 2 and E resulted in a false alarm with respect to the measure of the amount of the world searched with the friendly craft. These three test failures appear to be symptoms of one deficiency of the Crew E model.

The primary reason for these test results is that the Crew E model searched less of the world with the scout than did Crew E. This resulted in a greater amount of area searched exclusively by the friendly craft. This result is due to the fact that when both the scout and friendly craft have searched the same region, the measures used above credit the scout, rather than the friendly craft. The friendly craft measures reflect only the area that was searched exclusively by the friendly craft. Since Crew E searched more area with the scout than did the model, less area remained for Crew E to possibly search exclusively with the friendly craft.

The paths generated by Crew E and the model were consulted to identify why Crew E searched 11% more area with the scout than did the model. It was found that, in four of Crew E's seven sessions, the scout was flown in forested regions in the last few minutes of the session. The model provided a good fit to the area covered by Crew E up to the point at which the crew flew the scout directly into the forest. The scout path planning model does not generate paths through forested regions, as discussed in Chapter V.

One likely hypothesis that would explain this divergence between model and crew performance is that the crew was able to reevaluate the relative value of forested and open regions at the end of the mission whereas the model was not able to perform this reevaluation. At the point where the model and crew diverged, nearly all of the peak areas in the search affordance map (see Chapter V) were covered. Although the model includes a provision to avoid backtracking through previously searched regions, it still was not able to "realize" that, near the end of the mission, the only promising territory yet unsearched were the forested regions. The low affordance for these regions pertaining to difficulty in locomotion was able to keep the model from moving the scout into the forests. On the other hand, it is conceivable that Crew E flew the scout into the forests at the end of the mission because these were the only possible regions with any search value, regardless of the decreased speed of travel that would be incurred.

In the terms of the model constructs, Crew E was possibly able to redefine the search value vector that determined the relative worth of the different world regions. This vector determines which points in the world will have the highest affordance levels. The model did not produce this behavior because the search value was programmed in, and the model did not have the ability to "step back" and re-evaluate the degree to which this vector was an adequate representation of the search value of world regions. Although the theoretical possibility remains that Crew E's paths could be adequately matched by a suitably chosen fixed set of search affordance parameters, the attempt to find such parameters was unsuccessful.

The model could be given the ability for re-evaluation but this enhancement would be extremely computationally intensive. For example, instead of relying upon the static affordance maps that were created at the onset of the world display, the model could iteratively recalculate this entire map each time it "looked" at the map display. This recalculation would incorporate knowledge concerning areas that had been previously searched. At the end of the mission, the peak affordance areas in such a map would most probably be centered in forested regions, and the scout would therefore be commanded to these peak areas. Unfortunately, about five minutes of computer time are required to generate each map, so it would probably take tens of hours for such a model to perform one session given the current computer resources.

Conclusions

Five substantive deficiencies of the model were identified. With respect to the parameters used to mimic the behavior of Crew 1, it was found that the Crew 1 model may have been to rigid in its application of the serial cargo processing policy. It appeared as if Crew 1 was not as strict in the application of this policy. The serial processing policy is then best viewed only as a rough, but reasonable approximation of the behavior of Crew 1. This deficiency of the model resulted in the failure of four configural tests.

With respect to the deficiencies of the Crew 2 model, it was hypothesized that the assumption that all perceptual processors required equal time to be updated may not have been realistic. Rather, it seemed as if certain processors had to have been able to produce their affordances faster than could other processors. This model deficiency resulted in the failure of one similarity test and one configural test.

The Crew E model was hypothesized to be deficient with respect to its inability to dynamically reevaluate the search related world affordances. A potential enhancement to the model that would overcome this problem was discussed, although its implementation would be computationally intensive. This model deficiency resulted in the failure of two similarity tests and one configural test.

The use of a simplified set of friendly craft action commands caused differences between the models and the crews with respect to the average time spent by friendly craft above tree level. Unfortunately, no consistent policies for command selection could be identified that would allow a more realistic set of action commands to be employed in the model. This model deficiency resulted in the failure of three similarity tests and two configural tests.

Finally, the model could not be fairly compared with crew performance concerning idle time for the scout craft. The reason for this inability was that the idle time calculation required behavioral data at a level lower than the level at which the model generated behavior. This model deficiency resulted in the failure of two similarity tests.

In summary, the model failed on a total of 20 of the 132 tests. Sixteen of these 20 failures appear to be attributable to one of the major model deficiencies cited above. There is not guarantee, of course, that making the recommended enhancements to the model in an attempt to eliminate these failures would be successful, or that the enhanced models would not result in failures on a different set of performance measures.

These statistical tests have only been able to provide a measure of the empirical adequacy of the three parameterized models. Thus, the appropriateness of the unparameterized model, i.e., the model as a parametric family of specific models or as set of psychological assumptions, was not directly assessed. It would seem quite difficult, impossible perhaps, to try to directly evaluate the status of these assumptions, since any failures of a fully parameterized model based on them could possibly be due to an inappropriate translation of the assumptions into a testable model. This defense, of course, was not used above in an effort to defend the model, rather, the model failures provided some valuable and diagnostic information concerning the deficiencies of a few modeling assumptions.

Credit, though, might be given to the general modeling framework due to the fact that it allowed for a reasonably simple description of inter-crew differences in terms of the model parameters described in the previous chapter. The simplicity that is being suggested is not intended to be a reflection of the number of parameters that were required. Rather, the description of crew differences was economical due to the fact that a set of independent, relatively "low-level", parameters could be identified for each crew that were capable, via the model mechanisms, to generate the complex, "high-level", behavioral differences between the crews.

For example, it appears as if Crew E's ability, as a one-person crew, to perform at the level of Crew 2, a two-person crew, can in large part be accounted for by the simple fact that Crew E could fly the scout, enter friendly craft commands, and perceptually identify task affordances slightly faster than could Crew 2. These results may not be too surprising given the assumptions of the modeling approach discussed in Chapter IV. The offered view of skilled performance was that, once the appropriate perceptual and action mechanisms are in place, skilled performance even in complex tasks can be characterized as primarily as perceptual-response in nature. The "intelligence" in such processing was argued to be not so much a property of the human's real-time information processing, but rather as a property of the design of the mechanisms responsible for that processing. The fact that Crew E's ability to perform at the level of Crew 2 could be attributed to parameter differences entirely within the domains of perception and action is consistent with these views. Specifically, if skilled processing relies increasingly upon (albeit complex) perceptual and action mechanisms, the obvious place to look for performance limiting constraints is the perceptual and action domains.

If the performance limitations of the perceptual and action mechanisms cannot, on the other hand, account for all human performance limitations (as they could not for Crew 1), the next most obvious place to look is in the design of the perceptual and action mechanisms themselves. As opposed to the mechanisms' performance characteristics, their design characteristics do not reflect limitations on the rate with which the mechanisms operate. Rather, the design characteristics are concerned with the appropriateness of the higher-level strategies that the mechanisms can be described as implementing.

The most critical way in which design flaws in these mechanisms can yield sub-optimal performance is by their not being attuned to the appropriate set of concerns. Thus, the Crew 1 model was not only less efficient in the perceptual and action domains than the other two models, this model was attuned to a different set of world affordances than the other two models. The selection of action based on these affordances in the Crew 1 model was, in turn, less well aligned with the task goals than it was in the other two models. The assumptions of the modeling framework would suggest that the attempt to identify why Crew 1 might have been sub-optimal in this regard would have to take into account the historical forces in the previous task experience of Crew 1 that led to the development of the inappropriate perceptual and action mechanisms. Such an analysis is beyond the scope of the present work.

This Page Intentionally Left Blank

CHAPTER VIII

CONCLUSIONS

The research products selected to be presented here as conclusions are those with best chance of applying to skilled human performance in a wide range of task environments. These findings are intended to apply to skilled human behavior in perceptually rich domains that require the selection and execution of actions. The conclusions are also relevant to the assessment of alternative modeling approaches that were not specifically designed to describe cognition in such domains, but are nevertheless contended to serve as general purpose models of cognitive processing.

It is clear that the modeling framework used here falls far short of being an adequate theory of the organization of human perception, cognition, and action underlying skilled performance. This fact is true quite independently of any evidence provided for or against this framework in the empirical portions of this research. Due in part to its rough qualitative form, it is not even clear what data could serve to falsify its major assumptions, since any such failure could be due to an inappropriate translation of the qualitative claims into empirically testable hypotheses.

Even worse, its intended domain of applicability is so broad that the framework cannot be easily made to apply to any concrete task environment without careful study of the perceptually available information and the human's capabilities for action in that environment. On the other hand, the framework has provided a way of structuring the approach for investigating a rather complex phenomenon, and this purely pragmatic benefit is perhaps the research product with the best hope for applicability beyond the artificial laboratory task environment studied here.

One major assumption of the present approach is that much of the information processing work underlying skilled human performance can be performed by task-specific perceptual mechanisms that are sensitive to those aspects of the task environment that are highly relevant for the task of action selection. Using the terminology of ecological theories, the perceptual mechanisms have been described as mechanisms that are attuned to environmental affordances. These affordances are relationships between features of the task environment and the human's capabilities for action. Ecological terminology was used primarily to suggest that the assumed functional role of perception is to rapidly provide information that is high relevant to the task of action selection, given that the context in which perception occurs is compatible with this style of processing.

One glaring hole in the proposed framework is a lack of a specification of the properties of an environmental context that determine whether or not that context will support parallel processing. The clarification of this single issue is the subject of a large body of psychological research (e.g., Shiffrin and Schneider, 1977; Treisman and Gelade, 1980; Egeth and Jonides, 1972; Fisher, 1982). The research of Shiffrin and Schneider has focused on the primary role that consistency of training plays on the development of parallel processing, or automaticity. Treisman and Gelade's work, on the other hand, has focused primarily on the features of the stimulus display that either promote or disallow parallel processing. At this point, it is probably prudent to conclude that both training and contextual factors are relevant to the development of parallel processing.

To cloud the issue even further, the proposed framework can be interpreted in a way that does not even require perceptual processing to be parallel. The major requirement that has been made here is that the perceptual processing underlying expertise is rapid, and does not require multiple dimensions of information to be overtly and effortfully integrated. In some situations, this condition can be met without the assumption of parallelism. Rather, this type of rapid perceptual processing is compatible with the assumption that the perceptual systems mature by becoming attuned to the most diagnostic features of an information array.

Specifically, the perceptual mechanisms might be assumed to become sensitive to the minimal set of features that can be used to distinguish between relevant affordance-related situations. Two complex situations, even ones of very high dimensionality, may yet be distinguishable by their value along only one of the information dimensions. Under this interpretation, a complex situation is not recognized quickly due to the fact that a complex array of information is processed in parallel. Rather it is recognized quickly because it is the only relevant situation that has a particular value along one of its many information dimensions. Here, the focus would be placed on how an expert's perceptual systems can exploit environments of high redundancy, rather than how their perceptual systems might process large arrays of information in parallel. This latter interpretation is particularly attractive in that it would be consistent with the assumption that the amount of "raw" information processed by perceptual systems remains invariant over experience level, but the expert's perceptual systems produce more relevant information due to their sensitivity to only the most diagnostic information present in the visual environment.

If the proposed view of the role played by perception in expertise is subsequently found to be tenable, the following implications for psychological modeling would result.

First, the present view would encourage a holistic approach to the study of the perceptual, central, and action mechanisms underlying expertise. Before modeling central mechanisms, for example, careful attention would have to be paid to identifying the information that the perceptual system is capable of providing to presumably simplify the task of action selection. It is a widely held principle in artificial intelligence that finding a beneficial problem representation can simplify the processes that manipulate that representation to find solutions. In short, an "ounce" of representation is worth a "pound" of search. Perceptual systems can be described as the mechanisms that construct the problem representation used by the central systems. Ignorance of how the perceptual systems might simplify an information processing task by constructing an effective problem representation might result in models of central processing that are more complex than necessary.

Consideration, on the other hand, of how perceptual systems might construct easily usable representations may suggest simpler accounts of central mechanisms. This was the approach used here, and it allowed for a rather simple central mechanism to perform most action selection tasks, including planning and coordination tasks that seem to be suggestive of central processing of some complexity. In computer science terminology, the distinction between decision-making, planning, and coordination in the present model is a data, rather than process, distinction. That is, a unitary central process produces all three types of behavior; the type produced is determined by the data (the affordances) upon which the central mechanism operates. Perhaps the assumption that the model and human used the "same" type of central processing is not crucial to the present argument. Rather, what would seem to be important is that the modeling effort demonstrates that the assumption of simple central processing mechanisms in not inconsistent with the existence of complex human behavior.

Thus, the holistic approach would dictate that modeling of central processing should not proceed in ignorance of the fundamental role of perceptual mechanisms as representation builders. Similarly, the holistic approach would also suggest that modeling of perceptual processing should be influenced by those processes that exist "behind" perception. That is, instead of viewing perceptual systems as mechanisms that simply provide neutral descriptions of the environmental state of affairs, the role of action-oriented perception would be emphasized. While it may be useful (and perhaps necessary) to view some mechanisms of perception as being geared to providing a faithful picture of the external world, this research would suggest that perception can do much more. Rather than serving the benign role as a data gathering agent for the intelligent central systems, the perceptual systems can be viewed as intelligent mechanisms in their own right (Runeson, 1977). They operate intelligently by providing information that is highly relevant for the task of action selection, in a way analogous to an experienced secretary who might aid an executive by integrating knowledge of the executive's activities with the deluge of incoming messages, requests, and reports.

If found adequate, the proposed framework would therefore encourage an approach to investigating and modeling cognitive phenomena that views the perceptual, central, and action systems as tightly coupled. Note that this does not imply that these distinct systems cannot be decomposed for modeling purposes. But it has been suggested that the boundary lines that are drawn to reflect the decomposition may shift as a function of the experience level of the person being modeled. That is, an additional assumption of the proposed framework is that one of the primary psychological changes involved in skill development is a shift to a greater reliance on the perceptual systems, with a commensurate lessening of the burdens on the central systems.

It is interesting to observe that a shift to a greater reliance on perceptual systems would appear to have benefits associated with reduced memory demands. As discussed above, a neglect of perceptual systems when describing central systems can result in solution techniques that are overly complex. For example, Simon (1975) shows how the use of a perceptual strategy to perform the Tower of Hanoi problem can reduce the complexity of problem solving processes. In situations where possible perceptually-based solutions are ignored, though, the resulting "blind" models typically have another, yet related, undesirable feature. These models often require a large, often verbally represented, internal world representation (e.g., Winograd's (1972) SHRDLU program). A tremendous amount of logical reasoning and world knowledge are typically required in order to keep this representation internally consistent and externally valid (e.g., the "frame problem"). A sighted person, on the other hand, can look and see the effects of his actions, can see the progression of events due to other causal agents, and can thereby avoid much of the effort and many of the dangers associated with substituting a verbal representation of the world for the world itself.

The proposed approach, of course, is not intended to represent a general solution to the frame problem, or other problems associated with truth maintenance. Rather, what is suggested is that such problems might be largely irrelevant to the task of developing models of the cognitive processes characteristic of highly skilled behavior in perceptually rich task domains. Psychological processes might have a tendency to, where possible, substitute perceptual inference performed by searching the environment for logical inference performed by searching through verbal descriptions, or representations of the environment.

Due to the model's heavy reliance upon perceptual inference rather than logical inference, the current work encourages a view of representations as tightly coupled to the active processes of perception. This approach to specifying the nature of representations is in contrast to viewing representations as passive data structures that lie dormant waiting to be "matched" with features in the world. For example, the model's internal representation (the dynamic set of affordances) might be described as being "volatile". This term is meant to express the fact that the model's internal representation is dynamically reconstructed each time the model "looks" at the world display. The model does possess, of course, non-perceptual knowledge of world dynamics that it uses to create these affordances. This knowledge, though, is embodied in the constructive processes that generate the model's internal representation and should not be confused with the information that constitutes the model's representation at any point in time. This view of representations as objects that are dynamically constructed from perceptual information also serves to emphasize the important role the environment can serve as an external memory store.

To make these claims more concrete, perhaps it is best to see how the proposed framework might be applied for describing perceptual and cognitive processes in a more familiar domain. Empirical results concerning the behavior of expert chess players were cited in Chapter IV in order to motivate the present approach. One relevant finding was that chess experts were better able to construct meaningful board configurations from memory than were players with less chess proficiency. A meaningful configuration was one that might naturally occur in a chess match between skilled opponents. On the other hand, when configurations were comprised of pieces placed in arbitrary or unnatural positions, experts were found to be no better than the less experienced players at reconstructing these configurations from memory (Chase and Simon, 1973) The interpretation that was given for this result was that experts possessed superior pre-stored knowledge in the form of "chunked" chess positions that could be retrieved from long-term memory and rapidly deployed to encode the board position. Simon and Gilmartin (1973) have suggested that expert players may have accumulated about 50,000 such knowledge chunks.

Simon's interpretation of these empirical results does not, though, yield a natural account of DeGroot's (1965) observation that, for an expert, only a few candidate moves ever even "come to mind", and Dreyfus' (1979) observation that what sets an expert chess player apart is his ability to "zero-in" on the most important information. Simon might suggest, though, that some unspecified pattern matching process could be used to retrieve the appropriate chunks from long-term memory, and the activated chunks either are, or determine what, information comes to mind. Predicting what information the expert will zero-in on in a given situation could presumably be accomplished with knowledge of the set of chunks possessed by the expert.

A different interpretation of these findings can be given based on the modeling framework proposed in this research. This interpretation is incompatible with Simon's hypothesis that chess expertise is due in large part to the accumulation of chunked board configurations in long-term memory. Rather, the proposed interpretation is that much of the "knowledge" that chess experts have is literally in how they look at the board. This knowledge does not consist in stored patterns residing in long-term memory that must be matched with perceptual input, but in an array of chess-specific perceptual mechanisms that are sensitive to various features of the board configuration. Under this interpretation, chess knowledge does not passively reside as a set of representations, or data structures, within the chess expert. Rather, chess knowledge could be described as the functional ability of the player's chess-specific perceptual systems to dynamically construct highly relevant information from board positions, whether actual or internally imaged.

The perceptual mechanisms are assumed to process the board configuration rapidly, and, to some extent in parallel. These mechanisms are assumed to be attuned to the affordances in the board configuration. The affordances are the features of the board that are highly relevant for the selection of a move or series of moves. The player would presumably also be sensitive to his opponent's action affordances.

These hypotheses can be made more specific by reference to the model discussed in the previous chapters for describing behavior in the laboratory task. As in the perceptual mechanisms used for describing scout and friendly waypoint selection, it could be assumed that the chess expert uses a set of mechanisms that are each attuned to a different action-oriented feature of the task environment. For the waypoint selection model, these features were sensitive to world information relating to the search, locomotion, collision, and fuel-range affordances of world locations. For a chess model, the relevant affordances might relate to opportunities for center control, castling, mating, and the execution of various standard attacks and defenses. Each such mechanism might produce an independent mapping of the board that differentiates the board configuration to produce a scalar valued affordance map. As in the waypoint selection model, these maps could then be superimposed to identify moves or complex actions that have high affordance values on many of the affordance dimensions. There is nothing in the proposed framework that would suggest that it would be easy to rationally reconstruct the "computations" performed by these perceptual mechanisms in the format of algorithms.

To make the behavior of such a model consistent with the empirical claims made above concerning chess expertise, two additional assumptions would have to be introduced. The first concerns hypothesizing what it would mean for a move to "come to mind" in this model. A natural assumption would be to identify "coming to mind" with the process of passing the most highly afforded actions to the selection, or central, systems. In the waypoint selection model, the peaks and ridges in the final affordance map were passed to the selection mechanism for further processing. By a process of search and evaluation, one of these candidate waypoints was selected to be the next craft waypoint. This search process evaluated the candidate waypoints by introducing additional evaluative criteria to which the perceptual mechanisms were not sensitive. The evaluation of these additional criteria, concerning fuel maintenance and backtracking avoidance, could possibly have been performed in a perceptual mode by human crews, although perceptual mechanisms sensitive to these criteria could not be constructed in this research.

In a way similar to the waypoint selection model, the moves that could be described as coming to mind in the chess model would be those with the highest overall affordance value based on the outputs of perceptual processing. These moves would be submitted to the central systems to be analyzed in more detail to determine the appropriate course of action. This analysis would, as in the waypoint selection model, incorporate evaluative criteria to which the perceptual mechanisms were not sensitive, due either to insufficient maturity on the part of the perceptual mechanisms, or the fact that the evaluation of certain criteria is incompatible with a perceptual processing mode. In the case of immature perceptual mechanisms, the final affordance map may be "flatter", or not as differentiated as the map produced by mature mechanisms. Depending on an assumption concerning a threshold affordance value that determines which moves come to mind, the central mechanisms of a model incorporating such a flat map may either be flooded with options to be evaluated, or it may not "see" many appropriate moves.

Of course, another interpretation of unskilled chess behavior would be that the affordance map used by such players may be as differentiated as an expert's map, but the map does not capture the appropriate set of action affordances. For example, a less skilled player may not be sensitive to, for example, certain types of attacks or the value of center control. In such a case, the candidate set of moves submitted to the central systems for further review might be distinct from the candidate set of moves considered by the expert. A model based on this assumption would "focus-in", but it would focus-in on the wrong information.

To describe to superior ability of chess experts to reconstruct meaningful board positions, a memory-related assumption would be required. One possible assumption would be that board configurations are memorable to the extent that the affordance maps that are generated from them are highly differentiated or articulated. Perhaps the peaks or ridges in the map permit a chunking of information that enhances memorability. A less experienced chess player, with a flatter affordance map, would either lack or have less well articulated chunks than the expert. In the case of meaningless board configurations, neither the expert or non-expert would be expected to possess a highly differentiated map, since the normal patterns and "lines of force" to which the perceptual systems are sensitive would be lacking. Therefore, the expert's superiority in board reconstruction would be negated.

Hopefully, this application of the modeling framework to chess playing has helped to communicate the assumptions of the proposed approach. While this is far from being a complete model of human chess playing, it is hoped that this exercise has suggested that the modeling framework could be applied to describe perception and cognition in tasks that seem to require more than "simply" perceptual-motor skill. Perhaps more importantly, this application has helped to illustrate the vast array of issues left unaddressed within the proposed modeling framework, as assumptions had to be continually added to produce a model in agreement with the observable evidence. Whether the modeling framework can be enhanced to provide a substantial theory of expertise, even in restricted task environments, and whether such a theory could survive empirical test, are questions still to be addressed.

LIST OF REFERENCES

1. Allport, D.A., Patterns and actions: cognitive mechanisms are content specific. In G. Claxton, (Ed.), Cognitive Psychology: New Directions. Routledge and Kegan Paul, London, 1980.

2. Anderson, J.A. Cognitive and psychological computation with neural models. IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-13, 1983.

3. Baron, S., Zacharias, G., Muraldiharan, R., & Lancraft, R. PROCRU; a model for analyzing flight crew procedures in approach to landing. Proceedings of the 16th Annual Conference on Manual Control. Cambridge, MA, pp. 488 - 520, 1980.

4. Barsalou, L.W. Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. Journal of Experimental Psychology: Learning, Memory, and Cognition, Vol. 11, No. 4, 1985.

5. Chase, W.G., & Simon, H.A. Perception in chess. Cognitive Psychology, 4, 1973.

6. Chen, D. Hierarchical Plan Formation and Failure Recovery om the Flight Domain. Ph.D. Thesis, University of Illinois at Urbana-Champaign, 1983.

7. Churchland, P.S. Neurophilosophy. MIT Press, Cambridge MA, 1986.

8. Claxton, G. How do you tell a good cognitive theory when you see one? In G. Claxton (Ed.), Growth Points in Cognition, Routledge, N.Y., 1988.

9. Coulter, J. Rethinking Cognitive Theory. St. Martins Press, N.Y., 1983.

10. DeGroot, A.D. Thought and Choice in Chess. Moulton, The Hague, 1965.

11. Dreyfus, S.E. The risks! and benefits? of risk benefit analysis. International Journal of Management Science, Vol. 12, No. 4, 1984.

12. Egeth, H., & Jonides, J. Parallel processing in multielement displays. Cognitive Psychology, Vol. 3, 1972.

13. Ericsson, K.A., & Simon, H.A. Protocol Analysis: Verbal Reports as Data, MIT Press, Cambridge, MA, 1984.

14. Fisher, D.L. Limited-channel models of automatic detection: Capacity and scanning in visual search. Psychological Review, Vol. 89, No. 6, 1982.

15. Fodor, J. The Modularity of Mind. MIT Press, Cambridge MA, 1983.

16. Fodor, J.A., & Pylyshyn, Z. Connectionism and cognitive architecture: A critical analysis. Cognition, 28, 1988.

17. Gentner, D., & Stevens, A.L. (Eds.) Mental Models. Lawrence Erlbaum Associates, Hillsdale, NJ, 1983.

18. Gibson, J.J. The Senses Considered as Perceptual Systems. Houghton Mifflin Press, Boston, MA, 1966.

19. Harvey, N., & Greer, K. Actions: the mechanisms of motor control. In G. Claxton (Ed.), Cognitive Psychology New Directions, Routledge & Kegan Paul, London, 1980.

20. Hinton, G.E., & Anderson, J.A. (Eds.) Parallel Models of Associative Memory. Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

21. Hunter, J. On how we talk. In his Essays After Wittgenstein. University of Toronto Press, 1973.

22. Jacob, J.K., Facial representation of multi-variable data. In Wang, P.C.C. (Ed.), Graphical Representation of Multivariate Data. Academic Press, N.Y., 1978.

23. Jagacinski, R.J., Plamondon, B.P., & Miller, R.A. Describing movement control at two levels of abstraction. In P. Hancock (Ed.), Human Factors Psychology, 1987.

24. Johnson-Laird, P.N. Models of Deduction. In R.J. Falmagne (Ed.), Reasoning: Representation and Process in Children and Adults. Erlbaum, Hillsdale, NJ, 1975.

25. Kahneman, D., & Tversky, A. Prospect theory: an analysis of decisions under risk. Econometrica, 1979.

26. Keil, F.C. Mechanisms of cognitive development and the structure of knowledge. In R.J. Sternberg (Ed.), Mechansisms of Cognitive Development, W.H. Freeman Co., N.Y., 1984.

27. Kotovsky, K., Hayes, J.R., & Simon, H.A. Why are some problems hard? Evidence from the Tower of Hanoi. Cognitive Psychology, 17, 1985.

28. Lytton, L.E. A comparison of performance by one and two person crews in a simulated supervisory control task. Unpublished Masters Thesis, Department of Psychology, The Ohio State University, 1987.

29. MacKay, D.M. Evaluation: The missing link between condition and action. In W. Prinz & A.F. Sanders (Eds.) Cognition and Motor Processes, Springer Verlag, N.Y., 1984.

30. McHarg, I.L. Design with Nature. Natural History Press Garden City, N.Y., 1971.

31. Miller, G.A., Galanter, E., and Pribram. K.H. Plans and the Structure of Behavior. Holt, Rinehard, and Winston, New York, 1980.

32. Miller, R.A., Jagacinski, R.J., Plamondon, B.B., Lytton, L.E., & Kirlik, A.C. A comparison of one- and two-person crew performance in a supervisory control task. Proceedings of the Fourth Symposium on Aviation Psychology, Columbus, OH, 1987.

33. Murphy, E.F. (Ed.) The Crown Treasury of Relevant Quotations. Crown Publishers, N.Y., 1978.

34. Neisser, U. Preface to Concepts and Conceptual Development. U. Neisser (Ed.), Cambridge University Press, 1987.

35. Neisser, U. The ecological and intellectual basis of categorization. In above, 1987.

36. Neisser, U. Cognition and Reality. W.H. Freeman, 1976.

37. Neisser, U. Cognitive Psychology. Appleton-Century-Crofts, New York, 1967.

38. Newell, A. You can't play 20 questions with nature and win: projective comments on the papers of this symposium. In W.G. Chase (Ed.), Visual Information Processing. Academic Press, New York, 1973.

39. Newell, A., & Simon, H.A. Human Problem Solving. Prentice-Hall, Englewood Cliffs, NJ, 1972.

40. Nisbett, R.E., & Wilson, T.D. Telling more than we can know: Verbal reports on mental processes. Psychological Review, 84, 1977.

41. Norman, D.A. Categorization of action slips. Psychological Review, Vol 88, No. 1, 1981.

42. Norman, D.A. Reflections on cognition and parallel distributed processing. In J. McClelland & D. Rumelhart (Eds.) Parallel Distributed Processing, Vol. 2, MIT Press, Cambridge, MA, 1986.

43. Pew, R.W. A distributed processing view of human motor control. In W. Prinz & A.F. Sanders (Eds.), Cognition and Motor Processes, Springer Verlag, N.Y., 1984.

44. Plamondon, B.D. Dissertation research, untitled. The Ohio State University, Columbus, OH, 1985.

45. Pylyshyn, Z. Complexity and the study of artificial and human intelligence. In J. Haugeland (Ed.), Mind Design. MIT Press, Cambridge MA, 1982.

46. Raiffa, H. Decision Analysis. Addison-Wesley, Reading, MA, 1970.

47. Rasmussen, J. Skills, rules, and knowledge; Signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-13, No. 3, 1983.

48. Reitman, W., Nado, R., & Wilcox, B. Machine perception: What makes it so hard for computers to see? In C.W. Savage (Ed.), Perception and Cognition: Issues in the Foundations of Psychology. University of Minnesota Press, Minneapolis, MN, 1978.

49. Rosch, E.H. Cognitive representations of semantic categories. Journal of Experimental Psychology: General, Vol. 104, 1975.

50. Rosenbaum, D.A., & Saltzman, In W. Prinz & A.F. Sanders (Eds.), Cognition and Motor Processes, Springer Verlag, N.Y., 1984.

51. Rouse, W.B. Adaptive allocation of decision making responsibility between supervisor and computer. In T.B. Sheridan and G. Johannsen (Eds.), Monitoring Behavior and Supervisory Control. Plenum, New York, 1977.

52. Rouse, W.B., Rouse, S.H., and Pellegrino, S.J. A rule-based model of human problem solving performance in fault diagnosis tasks. IEEE Transactions on Systems, Man, and Cybernetics. Vol. SMC-10, 1980.

53. Runeson, S. On the possibility of 'smart' perceptual mechanisms. Scandanavian Journal of Psychology, 18, 1977.

54. Sheridan, T.B. Supervisory control. In G. Salvendy (Ed.), Handbook of Human Factors, 1987.

55. Sheridan, T.B. Toward a general model of supervisory control. In T.B Sheridan & G. Johanssen (Eds), Monitoring Behavior and Supervisory Control. Plenum, New York, 1976.

56. Sheridan, T.B. Three models of preview control. IEEE Transactions on Human Factors in Electronics, Vol. HFE-7, No. 2, 1966.

57. Shiffrin, R.M., & Schneider, W. Controlled and automatic processing II. Psychological Review, 84, (2), 1977.

58. Simon, H.A. The functional equivalence of problem solving skills. Cognitive Psychology, 7, 1975.

59. Simon, H.A., & Gilmartin, K. A simulation of memory for chess positions. Cognitive Psychology, Vol. 5, 1973.

60. Smolensky, P. Information processing in dynamical systems: Foundations of harmony theory. In Parallel Distributed Processing, Volume 1, Rumelhart, D.E., McClelland, J.L., and the PDP Research Group. MIT Press, Cambridge, MA, 1987.

61. Stich, S. From Folk Psychology to Cognitive Science. MIT Press, Cambridge, MA, 1983.

62. Treisman, A., & Gelade, G. A feature integration theory of attention. Cognitive Psychology, 12, 1980.

63. Tversky, A., & Kahneman, D. The framing of decisions and the psychology of choice. Science, 211, 1981.

64. Wason, P.C., & Johnson-Laird, P.N. Psychology of Reasoning: Structure and Content. Harvard University Press, Cambridge, MA, 1972.

65. Wesson, R.B. Planning in the world of the air traffic controller. IJCAI Proceedings, 1988.

66. Winograd, T. Understanding Natural Language. Academic Press, New York, 1972.

67. Winograd, T., & Flores, F. Understanding Computers and Cognition, Addison-Wesley, N.Y., 1986.

APPENDIX A

INSTRUCTIONS FOR SUBJECTS

I. Displays

A. Horizontal Map

1. Terrain - The horizontal map depicts a 100 square mile (10 mile x 10 mile) overview of the entire area accessible to the helicopters. If either ownship or any friendly craft exceeds the boundaries of this world, they will immediately be destroyed. There are three types of terrain:

   a. Open ground is represented as brown regions on the map.

   b. Lightly forested terrain is represented as light green regions on the map.

   c. Heavily forested terrain is represented as dark green regions on the map.

   The speed of the four friendly craft and the average speed of the ownship in automatic mode will progress from fastest to slowest in the following manner: above tree height or in open ground (fastest); in lightly forested terrain; in heavily forested terrain (slowest). The probability of being locked onto by enemy radar will also decrease in a similar manner from above tree height or open terrain (most vulnerable) to heavy forest (least vulnerable).

2. Home Base - Home base is represented as a white circle on the horizontal map. Upon reaching this region, friendly craft and the ownship may be unloaded, repaired, refueled, and resupplied with missiles.

3. Scout Helicopter or Ownship (F0) - The helicopter under your direct control is represented as a blue circle with the number zero. Your direction of motion is indicated by a small blue dot on the forward edge of the larger circle.

4. Friendly Craft (F1, F2, F3, F4) - Four friendly craft under your command are represented as four blue circles with the numbers 1 through 4. Their direction of motion is indicated in the same manner as for ownship.

5. Enemies - In the "world" in which you will operate there are at least four stationary enemies represented as yellow circles, at least four slow moving enemy tanks represented as orange circles, and at least four fast moving enemy helicopters represented as red circles. The direction of motion of the orange and red enemies is represented in the same manner as for the ownship. These enemies will not be visible on the horizontal map until they are detected by ownship or by friendly craft (see details below).

6. Cargo - There are at least eight pieces of cargo in the "world" that are represented as gray circles on the map display, each with a numerical identifier superimposed on it. The cargo will not appear on the map until they are each discovered by the scout or friendly radar (see below for details).

7. Crosshairs - A pair of crosshairs is represented as a black cross on the map display. The position of the cross is controlled by the two-dimensional joystick to the right of the horizontal map display. This cross is used to specify navigational waypoints for friendly craft and for ownship in the automatic horizontal control mode.

B. Text Display

This display is used to create and modify lists of goals and actions for the four friendly craft. (See below)

C. Forward Looking Display

This display presents a view of the terrain immediately in front of the ownship, with the maximum viewing distance equal to approximately 0.40 miles. Ground is represented in brown, trees are represented in green, and the sky is light blue. Enemies will appear as yellow, orange, or red circles. Cargo will appear as gray circles. Friendly craft will appear as blue circles. In the upper right corner of the display is a heading indicator, showing the direction (north, east, south, west) the ownship is pointed. The Forward Looking Display should be monitored so as to avoid collisions between the scout and either trees or friendly craft. Scout collision will result in its destruction and the termination of the session.

D. Status Display

The status display is immediately above the forward looking display and contains the following information:

1. Goal Lists - A list of the top (first) four goals and their associated actions for each friendly craft appears at the top of the display. These lists are modified by means of the text display and keyboard. (See below for details.)

2. Flight Management Information

   a. Fuel - Ownship and friendly ships will have a full load of fuel at the beginning of a mission. Rate of fuel expenditure will be increased by four multiplicative factors:

      1. Use of automatic horizontal control will increase the fuel expenditure rate by 100 % for the ownship.

      2. Going at higher velocities will increase the fuel expenditure rate up to 100 % for maximum velocity.

      3. Carrying greater cargo weight will increase the fuel expenditure rate up to 100 % for the maximum possible weight.

      4. Each time a craft is hit by enemy fire and damaged, the fuel expenditure rate increases by 25 % up to the maximum of 100 % for four hits.

      The combination of these factors can vary the rate of fuel consumption by a factor of 8 for friendly craft; 16 for the ownship.

      Fuel can be replenished by returning to home base. If the ownship or any friendly ship has less than one-quarter tank of fuel, a warning will appear in the message area, first when the one-quarter mark has been reached, then periodically after until either that ship's fuel has been replenished or the ship runs out of fuel entirely. If ownship or any friendly craft runs out of fuel away from home base, that craft is no longer operable, except for unloading its cargo.

b. Missiles - Ownship and friendly craft will have five missiles at the beginning of a mission. One missile is fired by a friendly craft during an attack (A) action, and one missile is fired by the ownship each time the trigger on the 3-dimensional control is pressed (see below for details). If a friendly craft is ordered to attack and it has no remaining missiles, the attack action will not take place, and a warning will appear in the message area. Missiles can be replenished by returning to home base.

c. Weight - The weight of ownship and each of the friendly craft depends on the number of remaining missiles and the amount of cargo that has been loaded into the ship. Each missile weighs 100 pounds. The combined weight of cargo and missiles cannot exceed 1,000 pounds. If a friendly ship or ownship tries to load a piece of cargo that would exceed this weight limitation, the load action will not be executed, and a warning will appear in the message area. Weight can be decreased by unloading cargo and by jettisoning missiles.

d. Damage Count - Each time ownship or friendly craft is hit by enemy fire, there is a 60 % chance the craft will be destroyed. If the craft is not destroyed, the damage count will increase by one. If the damage count reaches five, the craft will also be destroyed. Damage can be repaired by returning to home base.

e. Ground Speed - Ground speed ranges from 0 to 80 miles per hour in manual mode, and 0 to 60 miles per hour in automatic horizontal mode. The rate of fuel expenditure increases with ground speed. The maximum speed for any friendly craft will decrease by 10 miles per hour for each increment in its damage count.

f. Altitude - Altitude ranges from 0 to 200 feet. Zero altitude corresponds to ground level. In manual vertical mode, if the ownship reaches 0 altitude at a descending vertical velocity of greater than 25 feet/second, the ownship will crash and be destroyed. Tree level corresponds to 90 feet. Above this altitude, ownship and friendly craft can travel with the same ease as in the open ground area.

g. Vertical Rate - Vertical rate varies from -40 feet/second (descending) to +40 feet/second (ascending).

h. Mode Indicators - The ownship's presently operating horizontal control mode (manual, automatic without pathfinder, automatic with pathfinder) and presently operating vertical control mode (manual, automatic) are indicated. (See below for details).

i. Points - Your overall goal is to accumulate as many mission points as possible over each 30 minute mission. Points are awarded for delivering cargo to home base and destroying enemy craft. Points are deducted from current totals for losing any craft (including ownship) or for having a craft run out of fuel.

j. Time Remaining - Each mission will last 30 minutes. The remaining number of minutes are indicated on the status display.

3. Message Area - Warnings and error messages will appear in this area under the following conditions:

   a. Any friendly craft or ownship is locked onto by enemy radar. The ownship or friendly craft symbol on the horizontal display will also begin to blink on and off to indicate that it is locked onto.

      If ownship or a friendly craft is within .40 miles of an enemy, or, for the ownship, 1.5 miles while above tree level, then the probability of a lock-on during each one second interval is:

      1. 0.8 in open ground or above tree level

      2. 0.6 - 0.8 in light forest depending on altitude. The probability of a lock-on increases with altitude.

      3. 0.1 - 0.8 in dense forest, again depending on altitude.

      Once a lock-on occurs, the enemy craft will begin moving toward the ownship or friendly craft. If the enemy is not the currently active goal for the friendly craft that is locked onto (see below for explanation), the friendly craft will automatically begin an escape maneuver, and the goal on which the friendly craft was working will be suspended. If the enemy is the specified goal for the friendly craft, then the friendly ship will continue normal execution of the goal. Thus, if you want a friendly craft to approach an enemy without running away, the friendly craft must have the desired enemy specified as its current goal.

A lock-on may be broken in the following ways:

1. The distance between the enemy and ownship or friendly craft exceeds .40 miles (1.5 miles for the scout while above tree level). Since the enemies' lock-on range varies for the scout depending on whether it is above tree level (1.5 mile lock-on range) or below tree level (0.40 mile range), the scout, if above tree level, can break any lock-on of greater than 0.40 miles by descending to below tree level. If this does not break a lock-on, this means that the enemy is within 0.40 miles of the scout, and other action must be taken (fight or flee).

2. If ownship (locked onto at less than 0.4 miles) or a friendly craft transitions from open ground or from above tree level into any forested region, the probability of breaking the lock-on ranges from 0.2 to 0.9 depending on altitude (lower is better) and density of forest (heavy forest is better).

3. If ownship or a friendly craft transitions from light forest to heavy forest, there is also a 0.2 to 0.9 probability of breaking a lock-on, depending on altitude (lower is better).

If the lock-on is not broken within a variable period of time (mean is 14 seconds, S.D. is 2 seconds), the enemy will begin firing missiles at the friendly or ownship at intervals initially equalling the lock-on period, but decreasing slightly with each missile fired. The probability of being hit by each enemy missile depends on the density of trees and distance from the enemy. If ownship or a friendly craft is hit, there is a 0.6 probability of being destroyed. Otherwise, the hit will result in the damage count increasing by one.

b. Ownship or any friendly craft has less than one-quarter tank of fuel remaining.

c. A craft attempts to load a piece of cargo and any of the following conditions occur:

   1. The ownship is not at ground level, or a friendly is not "down".

   2. There is no cargo within 1/8 mile of the helicopter.

   3. The weight of the cargo will exceed the helicopter's carrying capacity.

   4. The scout is not stationary.

d. Any friendly craft is commanded to begin an attack action, and it has no remaining missiles (including ownship).

e. An invalid goal or action has been specified for a friendly craft.

f. A friendly or ownship attack resulted in the target being missed.

g. An invalid button sequence has been entered.

h. The pathfinder has failed to find a clear path from the scouts present position to the waypoint specified by a Go To.

i. The Auto/Go To button was pressed while the scout was moving too fast for the horizontal automatic mode with the pathfinder to be enabled.

All messages will be accompanied by an audio alarm to alert you to the fact that a message is being displayed. Messages remain displayed for approximately 10 seconds, then disappear. The disappearance of a message does not indicate that the condition that generated the message has been resolved.

II. Control of Ownship

A. Horizontal Control

The pilot will be able to control his craft in either manual or automatic mode. The advantage of automatic horizontal mode is that it frees the pilot from the manual task of tree avoidance. The disadvantage is that flight will generally be slower than in the manual mode due to dependence on an unsophisticated path finder, which is not as skillful as the human pilot. The automatic horizontal flight mode will also expend more fuel per mile traveled.

1. Automatic Mode (Row of three pushbuttons)

a. Auto/Go To Pushbutton - Automatic flight to a given location will be accomplished by moving the crosshairs to the desired location on the horizontal map display using the two-dimensional control stick, and the pushing the Auto/Go To pushbutton. A string of Go To commands can be entered by repeatedly positioning the crosshairs and pressing Auto/Go To. Each successive waypoint along the ownships' path will appear on the horizontal map display as a zero.

1. If the ownship is stationary when the first Auto/ Go To is entered, the helicopter will automatically proceed toward the postion indicated by the crosshairs, avoiding trees in its path. Path finding (tree avoidance) will involve only the necessary deviations from a straight line between ownship's current and specified locations.

2. If the ownship is moving at a ground speed of greater than approximately 20 mph when the first Auto/Go To is entered, the Go To point entered will not be accepted, the ship will come to a stop, and a message indicating that the scout was moving too fast for the pathfinder will be displayed in the message area. This is because the ownship must be stopped or moving very slowly for the pathfinder to be effectively engaged for the initial Go To of a string of such commands. Once the ship comes to a stop, a Go To can then be entered, or the ship can be returned to manual control, whichever is desired.

3. If the ownship is above tree height when an Auto/Go To is entered, the pathfinder will not be invoked and the ownship will fly in a straight line to the indicated point. If the ownship descends to below tree height during this flight, no automatic tree avoidance will be active as the craft continues along its route. If a string of Go To's is constructed in which the initial Go To was entered while above tree height, then the entire string of Go To's will be flown without benefit of the pathfinder, regardless of the altitude at which subsequent Go To's were entered. Thus, each Go To in the string will be treated as if it had been entered while the scout was above tree level.

4. When the craft finishes the last Auto/Go To, it will automatically stop. The ownship will also stop automatically for short periods if the pathfinder is slow in determining a path. There is no guarantee that a path can be found to any particular point given. If no path is found, a message indicating this fact will be displayed in the message area, and the pathfinder will turn itself off, erasing any string of Auto/Go To's entered after the point it was working towards. The above restrictions for pathfinder initiation make it very advisable to check the pathfinder status on the status display whenever the

   horizontal automatic mode is used.

b. Stop Pushbutton - Pushing this button will stop horizontal motion; manual rotation will still be possible. If the ownship is executing a string of Auto/Go To commands when the Stop Pushbutton is pressed, this action will erase that string.

c. Manual Pushbutton (Horizontal) - Pushing this button will activate the horizontal manual control stick, and will also erase any string of Auto/Go To commands that may have been entered.

2. Manual Mode

Manual Joystick (Horizontal) - The pilot will use a three-axis joystick as a rate controller. Rotation of the joystick will control the rate of rotation in the horizontal plane. The other two dimensions of the joystick control the direction and horizontal speed of the ownship. To move forward, backward, left or right, press the stick in the desired movement direction. The farther the stick is pushed from its upright resting position, the faster will be the the ownship's movement. Returning the stick to its upright resting position will cause the horizontal movement to stop. Also on this joystick are three buttons and a trigger. These are used for performing the following actions: left thumb button - LOAD, middle thumb button - UNLOAD, right thumb button - JETTISON, trigger - fire a missile. These actions will be described below.

B. Vertical Control

The pilot can control the vertical movement of the ownship in either an automatic or manual mode, regardless of which mode is being used for horizontal control.

1. Automatic Mode (Row of Three Pushbuttons)

a. Up Pushbutton - When this button is pressed, the ownship will automatically bob-up and maintain a standard altitude above tree height.

b. Down Pushbutton - When this button is pressed, the ownship will automatically bob-down and maintain a standard altitude below tree height.

c. Manual Pushbutton (Vertical) - Pushing this button will activate the one-dimensional vertical control stick for manual vertical control.

2. Manual Mode

Manual Control Stick (Vertical) - The pilot will use the one-dimensional control stick as a rate controller. The control stick will control the ownship's rate of change of altitude. Push forward to ascend, and pull back to descend. The farther the control is from its resting position, the faster the ownship will move. Care must be taken when descending, as reaching 0 altitude at 25 ft/sec or greater will result in the loss of the ownship and premature termination of the mission.

III. Control of Friendly Craft

To command the four friendly craft, it is necessary to specify a list of goals and the actions needed to accomplish each goal. This is done through the text display and keyboard which are activated by the Text Pushbutton on the navigator's control panel, or by the Interrupt Pushbutton on either the pilot's or navigator's control panel.

Text Pushbutton - Pushing any one of the friendly ship identifiers (F1-F4) followed by the Text Pushbutton activates the text display and keyboard. These are used to create a new goal list for the specified friendly craft or to modify a goal than has not already begun execution.

A. Plans of Action

1. Goals

a. E (Enemy) - One goal is to destroy enemies. There are a number of enemies scattered over the 100 square mile area accessible to the helicopters. These enemies will not be visible on the horizontal map display until the scout helicopter (ownship) is above tree level (> 90 feet) and comes within 1.5 miles of them, or until a friendly craft (including the scout when it is below tree level) comes within .40 miles of them. At these ranges, the scout or friendly craft radar will automatically discover the enemies, and display them on the horizontal map as yellow, orange, or red circles.

   1. Y (Yellow) - One class of enemies are stationary and are represented on the horizontal map display as yellow circles with a numerical identifier. Once discovered, these yellow enemies will remain displayed on the horizontal map until they are destroyed. When each is destroyed, it will disappear from the horizontal map, and mission points will be awarded. A friendly craft can always escape a yellow enemy's lock-on by fleeing since they are stationary.

2. O (Orange) - A second class of enemies are slow moving tanks, which are represented on the horizontal map display as orange circles with a numerical identifier. The direction of motion is indicated by a small orange dot on the forward edge of the larger circle. The displayed position of these orange enemies is accurate if they are within the radar range of the scout (1.5 miles above tree level, .40 miles below) or withing radar range of one of the friendly craft (.40 miles). If the enemies exceed these ranges, an approximate linear extrapolation of their movement will be displayed for 30 seconds. This extrapolated position will not be a true reflection of enemy craft location, but will be an estimate based on the craft's last observed direction and speed of motion. The fact that the enemy position displayed on the horizontal map is an extrapolation is indicated by the absence of the 'nose' or leading dot on the orange enemy symbol. This 'nose' appears whenever the enemy's position is being accurately displayed. After the 30 second extrapolation time, the orange symbol will disappear from the horizontal map. The orange symbol will also disappear if the enemy is destroyed, and mission points will be awarded. A friendly craft can usually escape the lock-on of an orange craft by fleeing since they are slower than the friendlies. Exceptions to this are when a friendly craft is 'down' in a forested region, or when a friendly craft has suffered 3 or 4 hits, in which case the friendly's speed will be less than or equal to that of the orange enemy.

3. R (Red) - A third class of enemies are fast moving helicopters that are represented as red circles with a numerical identifier. The description given above for the orange enemies also holds for the red enemies. The difference between the two classes of enemies is that the red enemies are faster and move more erratically than the orange. Thus, the extrapolation estimates of enemy craft position are less likely to be accurate for the red enemies than for the orange. Red enemies cannot be outrun by friendly craft since both have the same maximum speeds. As with yellow and orange enemies, however, lock-ons can be broken by the other means specified on pages 9 - 10.

b. C (Cargo) - Another goal is to pick up cargo and return them to home base. There are at least eight pieces of cargo scattered throughout the 100 square mile area accessible to the helicopters. Each piece of cargo will not be visible in the horizontal map display until the scout helicopter is above tree level and comes within 1.5 miles of it, or a friendly is above tree level and comes within 0.4 miles of it. At these ranges, 'radar' will automatically discover the cargo and it will be displayed on the horizontal map as a gray circle with a one or two digit numerical identifier. Cargo remain displayed until they are loaded onto one of the friendly helicopters or onto the scout itself. The cargo will vary in weight; however, the weight will not be known until the cargo is loaded. Mission points will be awarded for each piece of cargo that is delivered to home base.

c. X (Home Base) - Another goal is to return to home base to:

1. Drop off cargo
2. Replenish fuel and missiles
3. Repair damage to helicopters.

Home base is represented as a white circle on the horizontal map display.

d. S (Search) - Another goal is to have the friendly ship move about in an attempt to locate enemy craft and cargo. If a friendly craft comes within 0.40 miles of an enemy, the enemy will appear on the horizontal map. If a friendly is above tree level (up) and comes within 0.40 miles of a cargo, the cargo will appear on the horizontal map. Stationary yellow enemies will remain displayed until they are destroyed. Mobile orange and red enemies will remain displayed until the distance from the nearest friendly ship exceeds .40 miles, and the distance from the ownship exceeds 1.5 miles when above tree level, .40 miles when below. A linear extrapolation of enemy movement will then be displayed for 30 seconds, after which the enemy will disappear from the horizontal map.

2. Actions

a. A (Attack) - The indicated friendly craft ascends to above tree level, fires one missile, and descends to below tree level. If the friendly has no missiles prior to the attack, then the action will not be accomplished and a warning will appear in the message area. Friendlies and the scout must be within .4 miles of an enemy to have a chance of successfully attacking an enemy. Thus, for the scout, if it is above tree level and is locked onto by an enemy at a range of greater than 0.4 miles, any attempt to attack the enemy will result in a miss. If an enemy is visible in the forward-looking display of the scout, then it is close enough for the scout to have a chance to successfully attack it, although an enemy need not be visible to the scout (i.e. be in front of the scout) for it to be successfully attacked. Attacks can be successful, in which case the enemy attacked will be destroyed, or they can result in a miss, in which case the attacking friendly will still be in danger. The only means for increasing the chance of a successful attack is to get as close to the enemy as possible before initiating the attack. The probability of a successful attack by a friendly depends solely upon distance between the friendly and the enemy. As previously stated, the danger to the friendly of being attacked (fired at) by the enemy depends on the time since it was first locked onto by the enemy. The danger to the friendly of being destroyed once fired at depends on how far it is from the enemy and the type of terrain intervening between the friendly and the enemy.

   It is important to note that if, during an attack on an enemy, the attacking friendly misses the enemy with a missile, the missile can hit and destroy any friendly craft that may also be within range of the attacking friendly. So it is possible for one friendly to destroy another with a missile.

b. L (Load) - The indicated friendly craft loads cargo if three requirements are met:

   1. The cargo is within 1/8 mile of the helicopter. This will be true if the symbols for the friendly craft and the cargo overlap by more than one half on the map display.

2. The friendly is "down"

3. The additional cargo weight does not exceed the helicopter's 1000 pound carrying capacity.

If either of these conditions are not met, the load command will not be accomplished and a warning will appear in the message area. The goal string must either be fixed or aborted. (Reminder: In addition to these conditions, the scout must be stopped to load cargo.)

c. H (Hover) - The indicated friendly craft maintains current position and altitude. Hovers can be terminated with Skip, Interrupt, or Abort commands.

d. U (Unload) - This action causes the indicated friendly to unload all cargo that it is currently carrying. If the friendly or scout is not at home base when the action is executed, then the cargo become available for other craft to load, and no points are awarded. This could be issued, for example, to reduce weight or if the friendly has run out of fuel. If the friendly or scout is at home base, then mission points for the unloaded cargo will be awarded, and the craft will be repaired and replenished with fuel and missiles. Thus, the unload command serves to unload any cargo under all circumstances, and also serves to replenish and repair the craft if it is at home base.

e. P (Patrol) - The indicated friendly craft will begin a circular flight path. Patrols can be terminated with Skip, Interrupt, or Abort commands.

f. G (Go To) - The indicated friendly ship will travel in a straight line from its current position to the position indicated by the crosshairs on the horizontal map when the G button was pressed. The Go To position is indicated on the map display by a black number corresponding to the friendly number for whom the Go To was issued. A string of G commands can be entered by sequentially positioning the crosshairs at the desired waypoints, and pressing the G button.

g. J (Jettison) - The indicated friendly craft will jettison or eject one missile for each J command, provided it has one to jettison. This would serve to lighten the craft.

h. ^ (Up) - The indicated friendly craft goes to an altitude above tree level.

i. v (Down) - The indicated friendly craft goes to an altitude below tree level.

3. Format - The general format is to specify on each line of the text display a goal, followed by a colon, and then the desired actions separated by commas. In other words, the format of each line is

goal:action,action,action,etc.

a. E (Enemy) - To enter an enemy as a goal, enter E, followed by one of the three color symbols (Y, O, R), and then the appropriate digits to identify the particular enemy. These goal identifiers should then be followed by a colon and the action list.

b. C (Cargo) - To enter a cargo as a goal, enter C, followed by a one or two digit number to identify the particular cargo. These goal identifiers should then be followed by a colon and the action list.

c. S (Search) - To enter search as a goal, enter S, then a colon and the action list. Since the search goal is relatively unconstrained compared to the enemy, cargo and home (see below) goals, it is useful for inserting actions such as up or down into a goal list. For example, if a friendly craft is moving slowly through a forested region because it is "down", its goal could be interrupted and the goal

S:^

could be entered. After leaving the editor, the 'up' action would automatically be enabled since an interrupt was used (see below). It would execute, then the previous goal could then be resumed by enabling it.

d. X (Home Base) - To return to home base as a goal, enter X, then a colon, and then the action list. The X goal is different than the others in that it carries with it an implicit Go To which automatically brings the craft to home base. A usual action following the X goal is U. All actions entered in a home goal list will be executed after the friendly has returned home.

For example, commanding friendly craft F1 to attack red enemy number 6 might be accomplished as follows:

- On the navigator's panel press F1 followed by Text.
- On the text keyboard type ER6:G,G,A
  The crosshairs on the horizontal map would be carefully positioned before entering each G.
- Press Verify (See explanation below).
- If there are no syntax errors, press Exit (See explanation below). The goal and actions will then disappear from the text display and appear in the rectangle for F1 on the status display.

4. Editing Functions

In writing the goals and actions on the text display, the following editing functions will be useful:

a. Arrows (->, <-, ^, v) - Pressing each of these buttons moves the text cursor (the blinking rectangle) one space to the right, left, up, or down.

b. Open Line - Pressing this button sets up a new blank line at the cursor's present vertical position. This command is useful for entering a new goal into an already existing list of goals.

c. Delete Line - Pressing this button erases the line at the cursor's present vertical position.

d. Insert Character - Pressing this button puts a temporary filler character at the cursor's present position, shifting any characters to the right of the cursor over one space to the right.

e. Delete Character - Pressing this button erases the character at the cursor's present position, and shifts any characters to the right of the cursor one space to the left.

f. Next Line - Pressing this button moves the cursor to the beginning of the next line.

g. Verify - Pressing this button checks the line at the cursor's present vertical position for syntax errors such as a missing colon, missing comma, incorrect goal specification, etc. It also gives the coordinates of each Go To command in that line. Each new line should be verified in this manner.

h. Exit - Pressing this button terminates text editing and places the new list of goals in line for execution as indicated on the status display. No more than nine goals may be listed for a single friendly craft at any given time. Only the top four goals will appear on the status display. Whenever a goal is completed, it will be removed from the status display.

There are two additional buttons on the text keyboard - Reset and Repaint. Occasionally, the keyboard will lock during use and will not accept any input. To remedy this, the Reset button should be pressed. However, doing so causes the computer to output a message to the text display. This message can be removed and the screen restored to its pre-locked state by pressing the Repaint button.

A goal list can contain as many as nine goals of up to 18 characters each. The editor has a nine row by 18 column 'window' in which to enter goals and actions. Any attempt to move the cursor beyond this window will result in a 'beep' alarm being sounded from the text keyboard. This alarm will also sound if an attempt is made to enter a character on top of another, with the exception of the temporary filler character generated by the 'Insert Character' function.

B. Modifying Ongoing Action Plans

The pushbuttons mentioned next should be preceded by an identifier. The identifiers are:

F1 - Friendly ship 1
F2 - Friendly ship 2
F3 - Friendly ship 3
F4 - Friendly ship 4

1. Abort Pushbutton - Pressing this button cancels the currently active or top goal. The friendly ship then waits for an Enable to begin the next goal, if any.

2. Skip Pushbutton

a. If the friendly ship is performing some action when this button is pressed, the friendly ship will stop its current action and begin the next action for that goal. For example, pressing this button can terminate a Hover or a Patrol of the friendly ship.

b. If the friendly ship has just completed an action when the skip pushbutton is pressed, the friendly ship will skip the next action and begin the next action after that.

3. Enable Pushbutton - There are two uses of Enable:

   a. To initiate an attack.

   b. To begin a new goal.

   Prior to a and b the friendly ship will stop and wait until the Enable pushbutton is presse

4. Interrupt Pushbutton - At any time, either the pilot or navigator can interrupt a currently executing goal and the navigator can insert, via the text display and keyboard, a new goal/action list which will then become the currently executing action. The previous goal/action list will be placed second in line after the new goal/action list, and will consist of the goal identifier and any actions that had not been already completed at the time of the interrupt. Pressing the Interrupt Pushbutton performs the following:

   a. Suspends the currently active goal/action list.

   b. Turns on the text display, leaving the top line blank for the interrupt goal/action list.

   Pressing the Exit button on the text keyboard after an interrupt goal/action list has been entered performs the following:

   a. Turns off the text display.

   b. Enters the added goal/action list at the top the line for execution, moving all previous lists down one row.

   c. Enables the top goal/action list automatically.

Once it has begun execution, the interrupt goal/ action list can be controlled in the same manner as regular lists.

5. Status Pushbutton - Pushing this button will show the status of the identified friendly ship on the set of status indicators for approximately 10 seconds. After this time, the craft's information will automatically disappear and return to the ownship status. The status indicators display fuel remaining, altitude, vertical rate, ground speed, number of missiles remaining, weight capacity remaining, and damage count information that existed for the friendly ship at the time the Status Pushbutton was pressed. If, after pressing the Status Pushbutton for a friendly, ownship status is desired before the 10 seconds has elapsed, then the ownship status can be reinstated by pressing the F0 identifier followed by the Status Pushbutton. This is the only use for the F0 identifier button.