

AIDING HUMAN OPERATORS WITH STATE ESTIMATES

by

JAMES B. ROSEBOROUGH

B.S.M.E. Duke University
(1981)

M.S. Massachusetts Institute of Technology
(1984)

Submitted to the Department of Mechanical
Engineering in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1988

© Massachusetts Institute of Technology 1988

Signature of Author: _____
Department of Mechanical Engineering
May, 1988

Certified by: _____
Thomas. B. Sheridan
Thesis Supervisor

Accepted by: _____
Ain A. Sonin
Chairman, Departmental Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

MAY 25 1988

LIBRARIES
ARCHIVES

Abstract

AIDING HUMAN OPERATORS WITH STATE ESTIMATES

by

James B. Roseborough

Submitted to the Department of Mechanical Engineering on May 16, 1988, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Mechanical Engineering.

The problem of providing decision aids for operators of dynamic processes is considered. Specifically, the use of normatively derived state information including uncertainty as an aid to human decision making is studied. First, general issues regarding decision making and decision aiding are examined. Second, experiments are performed to identify some human characteristics in dynamic decision making tasks. Third, a decision aid is constructed for the retrieval of a tumbling satellite, and it is tested through experiments.

It is shown that objective evaluation of a decision aid may not be made in general. Thus decision aiding techniques can be evaluated only in the laboratory using simulated processes for which a model is known exactly. If a process model is available, the expected utility criterion may be used to determine a normative control procedure. The use of normatively derived state estimates appears to be a promising tool for decision aiding systems.

Experiments are performed to establish that normatively derived state estimates can be useful to human operators. Simplification models of the human are presented in which incoming information is simplified before modifying the decision maker's beliefs. In many cases human operators will make point estimate simplifications to presented uncertainty information. Rote procedure substitution, substituting a familiar process model, and other direct simplifications may also be made to reduce subjective workload.

A decision aid for retrieving a tumbling satellite in space is constructed and tested in the laboratory. Though a perfect aid improves decision making, an aid whose modeling errors are nearly equal to those of the human operator will degrade performance. In addition, an aid presenting only point estimates is superior to an aid that also presents uncertainties.

Thesis committee: Thomas B. Sheridan (Chairperson)
Derek Rowell
David L. Akin

Acknowledgements

I would like to thank my committee for seeing me through this research. First, I thank Dr. Thomas B. Sheridan for initiating and overseeing this project, and for bringing me into the MMSL. I also thank Derek Rowell for his help and interest in my work throughout my stay at M.I.T. David Akin deserves special thanks for coming to this project at a late date while already under a heavy schedule, and for sharing his expertise in space systems. Thanks also to Ed Smith who participated in the early stages of this project.

Many heartfelt thanks go to the members of the MMSL "lunch train" for their lasting friendship and support. In particular, I wish to thank Fo Buzan for regenerating my involvement in music, Chin Kan-Ping for his willingness to share and participate, Hari Das for his particularly well-timed conversation encouraging me to present, Max Mendel who helped me immensely with my presentation, Juggy Raju who has been a great friend in times of need, and honorary member Jon Simon who always has some interesting insights into everyday life.

The other members of the lab also deserve thanks for their help along the way, among other things for their often spirited conversations and their general interest in my well-being. This group includes Leon Charny, Cheng Chi-Cheng, Ren Jie, and Edie Greene.

I would like to extend sincere thanks to those people outside of the immediate lab who have provided me with an interesting and varied social life during my tenure at M.I.T. A partial list would have to include Dino DeAngelas, Lisa Deschènes, Julie Frost, Marsha Jeung, Joyce Lee, Tony Marra, Jayne Pierce, Ralph Vinciguerra, Steve Woolfson, and Yang Li-Li.

This research was conducted at the MIT Man-Machine Systems Laboratory. The Office of Naval Research has been very generous in supporting this work in its early, theoretical stages. Additional work presented herein was supported through the generosity of the NASA Jet Propulsion Laboratory. I wish to thank both of these sponsors for their financial support.

Table of Contents

Abstract	2
Acknowledgements	3
Table of Contents	4
1. Introduction and Thesis Overview	6
1.1 Decision Aiding of Human Operators	6
1.2 Considerations in the Design of Decision Aids	11
1.3 Descriptive Models of the Human Decision Maker	13
1.4 Design of a Satellite Retrieval Decision Aid	17
2. Decision Making for the Control of Dynamic Processes	21
2.1 Dynamic Decision Making Paradigm	21
2.2 Process Model	29
2.3 Normative Model of Decision Making	37
2.4 State Estimates as a Decision Aid	44
3. Considerations in the Design of Decision Aids	45
3.1 Statement of the Decision Aid Design Problem	45
3.2 Process and Task Characteristics	47
3.3 The Fixed and Partially Known Human Operator Characteristics	54
3.4 Decision Aid Characteristics	58
3.5 Systems Considerations in Decision Aid Design	64
4. Three Experiments in Human Decision Making	68
4.1 Introduction to Experiments	68
4.2 Experiment One: Three State Tracking Task	70
4.3 Experiment Two: Trapping Task with Negative Evidence of State	79
4.4 Experiment Three: Approximation to a Continuous Task	84
4.5 Simplification Models of the Human	90
4.6 Effects of Simplifications on an Ideal Decision Maker	97
5. Human Prediction of Naturally Occurring Rotational Motion	103
5.1 Experimental Design	104
5.2 Data Reduction Method	108
5.3 Preliminary Results	110
5.4 Major Results and Supporting Data	113
5.5 Implications for Decision Aids and Other Experiments	120

6. Decision Aid Design for Satellite Retrieval	123
6.1 Task Description	123
6.2 Task Analysis	125
6.3 Design of Decision Aid for Satellite Retrieval	128
6.4 Experiments with Decision Aid	136
7. Conclusions	146
7.1 Summary	146
7.2 Implications for Research in Related Areas	147
7.3 Recommendations for Additional Research	151
Appendix A. Vectors and Rotation Matrices	154
Appendix B. Quaternions as an Angular Position State Variable	156
Appendix C. Commonly Used Matrix Functions	158
Appendix D. Numerical Integration of Equations of Motion for a Rotating Rigid Body	161
Appendix E. Estimator System Equations Definition	163
Appendix F. Bayesian Estimation of Model Parameters	165
Bibliography	168

1. Introduction and Thesis Overview

This chapter, consisting of four sections, briefly introduces the material contained in the thesis. First, some general issues are presented and the scope of the study is given. Then, a qualitative analysis of the decision aid design problem is introduced. Third, simplification models of the human decision maker developed in this work are discussed. Finally, the results are applied to the problem of retrieving a tumbling satellite in space.

1.1 Decision Aiding of Human Operators

1.1.1 Scope of Systems Studied

We are considering the problems of providing decision aids for systems in which there is a human operator controlling a dynamic process. Additionally, we are interested in systems where the human operator is considered a necessary part and plays an active role in the decision making. There are many examples including power plants, airplanes, ships, process control plants, space missions,

the control of economic systems, political systems, military systems, and others.

1.1.2 Reasons for Man-In-The-Loop

We present four major reasons that a human operator is considered necessary in a control system:

- Model of controlled process is inadequate. The human is included because of his extraordinary judgement and reasoning powers when the process attains some unmodeled state. Examples to which this applies include space, power plants, process control plants, economic, and military systems, typically because of the difficulty in modeling or inability to anticipate the exact nature of the environment.
- Best alternative cannot be computed. Multiattribute decision problems [10] are by their nature incomputable and can arise in a variety of domains including politics and management. Computational limitations because of problem size can also cause this problem.
- Human's unique sensory abilities cannot be replaced. An airplane pilot's ability to see the airport or a power plant operator's ability to relate sounds to specific conditions are good reasons to have a human present in these systems.
- Political factors prevent replacement of human. For example, engineers find it easier to put responsibility on a fallible operator than take responsibility themselves. Or, putting astronauts in space receives more media attention than unmanned space missions, so may result in more project funding.

1.1.3 Desire To Provide Decision Aids for Operators

There are several reasons, both bad and good, to expect interest in decision aids presently and in the future. Recently, people have begun to realize the enormous potential of computers, and are looking to the computer as a panacea for any frustration, including the frustrating losses caused by man-related disasters. In response to a disaster, the engineers diagnose the complex man-machine system as having a specific problem that needs to be fixed. Loosely speaking, this leads to a patch solution by the engineers. The logical patch in a man-machine system is the addition of a decision aid. The flexibility of the computer gives the frustrated public hope that decision aids are capable of preventing disasters.

Note that in purely mechanical domains, this type of problem-cause-fix behavior generates successful designs. A carburetor may be invented that operates successfully in nominal conditions. When it behaves poorly in cold weather, however, the problem is diagnosed as varying gasoline evaporation rate and a choke mechanism is added as a fix. The same approach has poor side effects in man-machine systems. While the independent pieces of a mechanism operate simultaneously to produce coordinated physical behavior such as car motion, the separate pieces of information presented to a human operator cannot be operated on simultaneously by the human to produce coordinated, complete information processing. Thus as systems become more complex, the resource-limited human operator must become skilled at ignoring those pieces of information that are not important to the task.

1.1.4 Expecting Continued Interest in Decision Aiding

The reasons we can expect continued interest in decision aids are the same reasons we can expect the human to remain in the loop of many dynamic systems. As technical processes and society get larger and more complex, there will always be systems and problems that are at the limits of our understanding or ability to control. We will then ask human operators to use their singularly intelligent abilities to control these systems. Since historically the public has been skeptical of new technology, we can expect them to mistrust future technologies and demand, as they have in the past, that humans rather than machines take responsibility for the consequences. They will continue to demand that a human be in control, perhaps even after the human can be outperformed by an automated system. There is no reason at present to expect this mistrust of new technology to change on a significant scale.

1.1.5 Problem of Assessing the Value of a Decision Aid

Suppose we accept that a decision aid may be useful in a given context, and we generate a design for one. Normally, we would measure decision performance first with and then without the aid to determine its value to the system. However, for a large, significant group of decision aiding problems that are not primarily motivated by political factors, the positive value of a decision aid is difficult or impossible to establish by objective means.

To see this, note that for a given process control situation, either the best possible action can be computed or it cannot be computed. If the normative

action can be computed, then the normative action should be used in place of the human decision and the human operator is not necessary. Since we are considering systems where the human operator is necessary, it follows that the normative action cannot be computed, and a single decision cannot be judged as to its value, optimality, or correctness. Thus the decision performance must be evaluated by some other means.

The usual way to attempt this is to measure the average decision performance with and without the decision aid. Again, decision aids are typically desired for systems involving high risk, where the cost of certain mistakes is high compared with the costs incurred during normal control operations. Measuring average decision performance requires seeing several events, and for high risk, low frequency events, this is simply impractical; we cannot wait for several nuclear accidents to determine whether they were preventable. Simulation in the laboratory is a partial answer to this problem, but the low frequency problem remains. In addition, the simulation must be of appropriately high fidelity, perhaps high enough to be able to determine normative control actions.

1.1.6 Subjective Methods of Inquiry

If objective measures of evaluation cannot be determined, Then the designer must attempt to understand as well as possible all elements of the system and their interactions. This means understanding all aspects of the controlled process, the human operator, and the decision aiding system.

Each element may be studied in the laboratory using completely modeled

processes. While the laboratory process is a limited version of the real domain, the experimenter may retain a complete, exact description to be used in evaluating laboratory generated decisions. In this work, we have experimentally examined human decision making in controlled laboratory conditions to develop some descriptive models of human decision making behavior. The problem of retrieving a tumbling satellite was chosen as an example decision problem whose simple elements can be brought into the laboratory and examined carefully.

1.2 Considerations in the Design of Decision Aids

1.2.1 The Decision Aid Design Problem and Analysis

Like other design problems, the decision aid design problem consists of selecting a good set of characteristics for the decision aid from those available. Of course some designs are better than others in that the average system output will be higher with one than the other. It is the role of a qualitative analysis to predict which designs will be better than others. As we have already noted, this subjective approach may be the only available approach for some important systems.

The system to which a decision aid will be added consists of a controlled process, a human operator, and a decision aid. Each of these components has characteristics associated with it. The designer considers the characteristics of the process and human to be fixed, and those of the decision aid to be selectable. The designer must select decision aid characteristics which he feels will maximize system performance. Examples of process characteristics include process rate, the degree to which the process is modeled, the degree to which the process is observ-

able, the degree to which the process is predictable, complexity, and the structure of the utility function. These characteristics serve to differentiate various problem types, and should be considered in a qualitative analysis.

The decision aid will have its own set of characteristics. A characteristic that is considered in detail in this work is the use of probability information and process state as part of the decision aid. The fixed and partially known characteristics of the human operator are important to determine beforehand, to the extent that they can be. In our example decision aid problem, we first measure the abilities of the human in observing and predicting rotational motion.

Each set of process characteristics will influence what are desirable and undesirable decision aid characteristics, and ideally this is fully predictable and insensitive to the detailed characteristics of the process. For example, a flexible decision aid will be useful primarily when no explicit process model has been determined beforehand, and will be of less use when there are severe time constraints in the task. Chapter 3 examines some of these interactions.

1.2.2 Systems Considerations in Decision Aid Design

Taking a systems point of view means considering the process-operator-aid system as a whole rather than as a collection of independent parts having individual characteristics. Changing one part within a complex system often has many unanticipated effects. By taking a global, systems viewpoint, more of these effects are anticipated. This allows the average positive or negative effects of a potential decision aid to be more accurately predicted.

In the systems view, the complex system is the unaided process plus the human operator, and the change to the complex system is the addition of a decision aid. This change could degrade system performance if the increased workload overwhelms the operator, for example. Or, the operator may rely on the aid to such a degree that he cannot function without it, a situation which may be disastrous should the decision aid fail at a critical time.

1.3 Descriptive Models of the Human Decision Maker

We now briefly present some descriptive models of the human decision maker, both in the literature and from experiments.

1.3.1 Models from the Literature

A common model of the human decision maker is the sequential processor having limited resources [7][17][21][31][32][37][55]. In this model the human is a sequential computer limited in his processing rate and time-shares some of his resources to attend to multiple tasks simultaneously. The human has a vast long term memory, but is limited to several short term memory elements. However, some parts of human information processing are most certainly parallel. Examples are most of the visual system, or attention mechanisms that alert us when we are engaged in conversation but our name is called elsewhere.

We shall combine the models of the human as a mistake generator [48], or imperfect calculator. There is a large body of literature devoted to human error and its causes, but when we talk about decision aids, we are generally talking about more than preventing mental slips. It is well established that people do not

update their beliefs from new evidence in a Bayesian way [16]. Nor do they work well with small probabilities; the number of people playing state run lotteries will attest to this. There are other predictable human biases that lead us to call the human an imperfect probability calculator.

Rasmussen [44] has proposed a knowledge, rule, and skill based hierarchy of cognitive decision behavior. Roughly, it says that the novice learns a task through the application of rules that have been given to him in training. When many of these rules are internalized, the decision maker has acquired skill-based behavior, which is the intermediate level of proficiency. Finally, an expert operator has on hand an elaborate knowledge base that may be used to determine responses to new circumstances. Knowledge based behavior allows the human to make good decisions when the process is in an unfamiliar state.

There is at present a great interest in the so called "mental model" of the controlled process [13][26][28][35][36][69], which the operator is continually accessing to predict the outcomes of his actions and finally select one. Conant and Ashby [12] have written that any successful controller of a process will contain a model of the process. This latter use of the word model refers to an approximate correspondence between the set of process states and a set of mental states. "Mental state" here is broad; in a rule based operator it could be interpreted as the present understanding of the applicable rules. As we allow "mental model" to have such broad meaning, we let the term "decision aid" encompass nearly all equipment besides the controlled process and the human operator.

In the area of man-machine systems, there is also a history of studies into human tracking abilities [4][25][30][39][40][63][64]. The concept of the internal model concept has received more attention recently [19][24][56][65][68], and there are presently many general modeling efforts directed at the supervisory control problem [6][8][18][23][41][43][45][50][51][54][62][66][59].

Finally, classical decision theory concerns itself with the concept of a normative decision maker who always selects the best action from those possible. For the normative decision maker, the belief state and utility function are separable components. The human may be modeled as an imperfect executer of the normative decision making algorithm. In the next section, we discuss some specific imperfections in the normative algorithm that produce successful descriptive models of human decision making in experimental tasks.

1.3.2 Simplification Models of the Operator

As the result of experiments that are reported in Chapter 4, a few specific models of sub-normative behavior in human decision makers are proposed. These models together are called simplification models because in each case the decision maker appears to use a mental algorithm that is simpler than the complete normative decision making algorithm. Often this is done with only slight effect on the overall decision performance. In some sense all engineering modeling consists of making simplifications that retain most of the use and predictability of the original system, and it is in this spirit that we model operators as making these simplifications.

The ideal decision maker is considered to have a belief state about the state of the controlled process, and a utility function that represents the task goals. To maintain his belief state, the operator must be able to predict future process states from his present belief state and interpret observable indications of state. Roughly speaking, these require models of the process and models of the observables. Finally, he needs a way to combine his beliefs and utility function to select a specific action.

Apparently to reduce mental effort, the decision maker may use simplified forms of the corresponding normative elements. Three of these possibilities have been observed in experiments. They are,

- The human may substitute a memorized rote procedure for the complete normative procedure. He responds to observables in a predetermined way, possibly by operant conditioning. For certain types of tasks, this may not have a large effect on his average decision making performance.
- The human may extract certain parts of presented state information and ignore the rest. The phrases "partitioning," "making point estimates," and "focusing" will be used with this type of simplification.
- The decision maker may use a familiar inaccurate model of the process or observables in place of an accurate model. As if to avoid the effort of learning new process behavior, decisions are made as they would be for a similar, well understood process.

1.4 Design of a Satellite Retrieval Decision Aid

To examine human decision making, a decision aid was constructed for retrieving a tumbling satellite. We will highlight the results of these investigations.

1.4.1 Statement of Satellite Retrieval Problem

In space applications, it will become more common to service previously launched artificial satellites. These satellites will have some latent angular momentum, and some may be undergoing tumbling motion. When the satellite has a tumbling motion, it is difficult for the human to predict future positions or states of the satellite. The problem is to design a decision aid for the retrieval of satellites having tumbling motions.

The previous approach to retrieving satellites involves flying around the rotating satellite at the same rate as the spinning object, then closing in on the object [11][22]. Though this has certain advantages such as reducing interaction forces on reaching the satellite, controlling the vehicle can be difficult for the human when the target has simple spinning motion, and extremely difficult when it has tumbling motion.

As an alternative, we suggest that the human first achieve a parking orbit adjacent to the tumbling object. Then he may observe it for a period using automatic equipment such as the decision aiding system until a model of the tumbling motion is built up on the computer. When this has been done, the decision aid may display information about future states of the satellite, and a robotic arm

may be used to grab the fixture when a good opportunity arises. The decision problem consists of determining when a good opportunity has arisen.

1.4.2 Satellite Retrieval Task Characteristics

The controlled process here is the tumbling satellite, in the absence of forces or torques. A primary feature of this task is that, while the motion of the satellite is extremely predictable given its present state, human operators have difficulty making such predictions. The state is observable for a computer, and the process is capable of being well modeled.

The task is a one-time decision, with time pressure of little or no importance. The operator must judge when a good opportunity exists to grab the satellite. The task has a clear target state, to be grappling the satellite. The task is unforgiving, since an attempted grapple either results in some form of success, all of which are equal, or failure, which can worsen the spin of the satellite and send it elsewhere. Finally, there is little that is complex or hierarchical about this task as the effects of the possible actions are obvious and easy to describe.

1.4.3 Decision Aid Design

Because prediction is important for good decisions in this task, the decision aid must be able to measure the satellite's state, predict its future states, and convey these to the operator. A decision aid was designed that includes predictions based on present estimates. In addition, historical traces of past estimates can be examined to establish a base level statistical model of the process.

To respond to the natural hierarchy in the task, a "grappleability" function was defined over the process state space that measures how far the present state is from the most desirable state. In addition, an acceptability measure relates values of the grappleability function to a probability of task success. Experiments were performed to determine which forms of information are most useful to the human operator.

1.4.4 Experimental Results

In the experiments, it was established that decision aid based on these concepts was helpful. As expected, as the fidelity of the decision aid was increased, decision making performance improved.

An important point is that when the human operator was given a poor decision aid. Here poor was defined as having modeling errors nearly equal to those of the human. For this case decision making was reduced to a level below that of decisions made with no aid. The mechanism behind this was apparently the attention sharing required to combine direct visual information with the decision aid. When a good aid was available, direct visual information was not used; the human operator directly implemented the aid's recommendation.

Another result is that the performance measures that are improved are related to the information selected for use by the operator. Since subjects had a choice of information types, a single subject naturally settled on using a single source, possibly differing from that of other subjects. The performance measure that was improved had a strong relationship to the information used. Turned

around, if a task consists of maximizing points received, then a display showing expected points for possible actions will produce the best performance. An implication is that if an operator is given a choice among informational displays, he may choose to rely on the wrong display since only one will result in superior decision making performance. From both of these results we may say that where information is concerned, more is not always better.

In addition, performance was improved when only point estimates were presented, compared to performance when uncertainty was also presented. Hence an overconfident decision aid appears to be better than one which is honest about its uncertainty. Apparently this is because when presented with uncertainty data, people will only use point estimates which they must work to extract themselves. While nature abhors a vacuum, humans abhor uncertainty.

2. Decision Making for the Control of Dynamic Processes

In this chapter we present a general process model called the partially observable Markov model. This model can represent a wide variety of processes, and will be the basis for quantitative discussions of many aspects of decision making. We begin by introducing the system consisting of the process plus decision maker plus decision aid. The analysis in this chapter serves as a basis for the numerical analysis of later experiments and analysis reported throughout the remainder of the thesis.

2.1 Dynamic Decision Making Paradigm

2.1.1 Separation of Process and Decision Maker

The basic elements of the man-machine system before a decision aid is added can be separated into the controlled process and the human operator or

decision maker. This is shown in Figure 2.1. Coming from the decision maker and going to the process are inputs, or actions, denoted $u(t)$. Coming from the process and used by the decision maker are the process observables, or outputs, denoted $y(t)$.




Figure 2.1. The decision maker uses process observables to select inputs to the process.

When a decision aid is added, the system is as shown in Figure 2.2. From the operator's viewpoint, there are the normal system inputs and outputs, and additional inputs and outputs of the decision aid. In addition, the decision aid may directly use the process observables and directly manipulate system inputs. Many things can be a decision aid according to this view. The decision aid may simply execute the low level control loops that the human operator has requested, or it may do more, such as suggest appropriate actions to the human.

2.1.2 Action-State-Observable-Belief Model

An essential element of the dynamic process is a state $x(t)$ which is distinct from both the inputs and the observables. This is depicted in Figure 2.3. To observe and control the process is to observe and control this state.

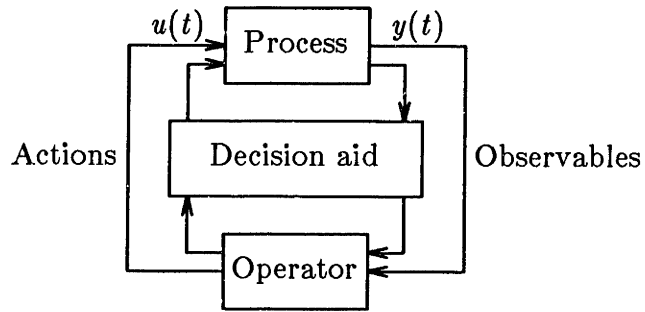


Figure 2.2. A decision aid communicates with both the process and the human decision maker.

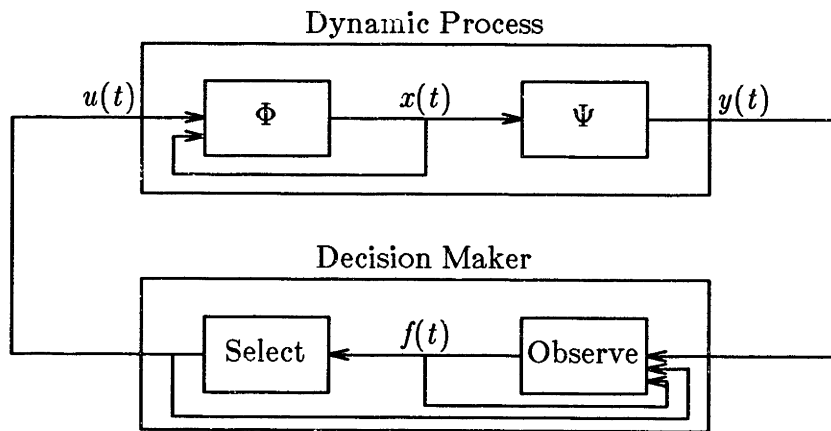


Figure 2.3. The process has a state not available for direct observation, the operator has beliefs about this state collectively called the belief state.

For an ideal human controller, we assume he has a set of beliefs about the state of the process. We call this set of beliefs the belief state $f(t)$. As decision theorists, we are interested in the time dependent relationships among process state, observables, belief states, and actions. If we could predict the interactions among these elements, then it would be possible to predict the system behavior as a whole. In practice we cannot predict the interactions to such a degree, but it is partly the goal of this work to describe them as accurately as possible.

A more detailed model of the human elements would consider separate components for the perception of the observables and the carrying out of intended actions. Since we are interested in the deep cognitive behavior of the human, we assume that perception of displayed values and execution of intended actions are perfect in the human. Often such an assumption is justified if the interface is well designed and time constraints are not too heavy. At other times, we will have to consider the degradation of system performance because of limitations in these additional human elements, however. We take this up again in Section 2.1.4 and elsewhere.

2.1.3 Necessity of Action and Decision Making

In this work, as in classical decision theory, we take the view that when one takes no action, one has acted by making the decision to take the null action. A dynamic process will continue to march through its succession of states whether or not the human operator chooses to take specific and new actions or change set-points. It is also in some contrast to the way decision making is discussed commonly, where an operator considers the information before him, weighs it carefully, and finally comes out with "a decision." In that view there is no decision and the universe is in limbo until the decision is explicitly made by the decision maker.

In our view, it may be valid to model a real decision maker as going through decision making activities in a sequential process, but all the while he is doing this he is unwittingly deciding to "leave all inputs alone" until he feels his mental processing of information is complete. Thus a record of the decision made

at an instant consists of the values of each setpoint, switch, and so on at that instant. For the most part we will consider discrete time tasks such that the human has a short, specific amount of time to think between times when the decision is recorded.

2.1.4 Form, Content, and Coding of Information

We now consider a more detailed picture of the two interfaces between man and machine, the action interface and the observable interface. Figure 2.4 depicts the basic elements of these interfaces.




Figure 2.4. Action and perception in the man-machine interface.

In this view, input devices or sensors are used for the purposes of measuring the intention of the operator, while display systems are used to convey specific information to the operator. The real world variables that transfer the information may have high dimensionality, may be difficult to describe analytically, and are not unique in a given context. The success of an interface may be characterized by the correspondence between human information and process information.

In the ideal case, all the process blocks in Figure 2.4 are perfect transformers of information from one representation system into another. Then the meas-

ured intention is equivalent to the intention of the operator and the process observables are perfectly understood by him. In practice the sensor and display elements may come close to being ideal translators, while the human's perception and execution are far from ideal. In considering whether to replace specific elements of the human's decision algorithm with the computer, one must keep in mind the additional interfacing requirements and the misinformation that will arise. In addition, an interface may require conscious thought by the operator, thereby increasing his workload.

2.1.5 Hierarchy in Actions, States, and Observations

In much human activity, the execution of actions and the perception of observables have strong hierarchical elements. For a human, actions can be composed of subactions, and observables can be composed of subobservables. As examples, consider the action of beginning laundry:

Beginning one's laundry:
Open lid.
Put clothes in washer.
Add 1/2 cup laundry detergent.
Close lid.
Turn selection knob to "Colors."
Insert money.

or the following example of an observation:

If you see
1) A very bright star
2) with a tail behind it
3) which moves each day
you have seen a comet.

Hierarchy enters into the dynamic control environment as well. Here we must add that there is also hierarchy in the conception of process states. Examples of this type of hierarchical construction would be composite states such as, "Normal State," or "State of Emergency." What is meant by "Normal State" is really any of a large number of basic process substates.

2.1.6 Mixture among Hierarchical Elements

Far more common than the simple state-substates examples of hierarchy from the previous section are hierarchical concepts that are mixtures of actions, states, and outputs. Consider the following example, which is a first level expansion of the action, "To go to Bill's house."

To go to Bill's house:
Go until the third light.
If you *see a church*,
you have gone too far.
Turn left
...

Note that this action consists of subactions, **substates**, and *subobservables*. If one follows the directions, he should never see a church, but if he does see a church, his next action is indeterminate. A person carrying out the action and seeing a church would turn around and eventually find Bill's house. Curiously, the action definition is simultaneously incomplete and robust to mistakes in carrying out subactions. Examples can also be given of high level observations that may require actions to test, or preconditions on process states to conclude. Such a definition is often called a procedure or algorithm, although use of the latter implies that a specific outcome is guaranteed, a condition that is not met in the

"Bill's house" example.

Another common example of hierarchical action in dynamic control systems is the decision to execute a control law of the form $\mathbf{u}(t) = \mathbf{K} \mathbf{y}(t)$. The entity that carries out this action will of course be continuously making measurements $\mathbf{y}(t)$ and taking subactions $\mathbf{u}(t)$, but the decision maker will be taking a single action, "implementing linear control law." Generally speaking, the actions at higher levels of the hierarchy are changed less frequently and are therefore preferable for human decision makers. This is the basis for supervisory control [52][53].

2.1.7 Deterministic and Stochastic Elements

We take the view that underlying all physical processes is a deterministic process operating at some fine level of detail, the true behavior of which we can never know. That we can never know an exact deterministic model of a physical process should not stop us from building models of processes. In any real, practical modeling effort, we eventually stop modeling the fine detail of the process behavior at some point and lump the rest of the process behavior into "noise." This is a simple picture of the basic distinction between deterministic and stochastic elements within a system model.

In the sections to come, stochastic models will figure heavily in the analysis and experimentation. Initial stochastic models will be combined or separated to form new stochastic models that have more or less correspondence to the real, underlying deterministic system that is being modeled by both man and machine. Stochastic elements are used by choice, and allow us to stop our search for the

deterministic truth and start doing something with what we have modeled so far.

2.2 Process Model

We now present a process model that will serve as the basis for experiments in Chapter 4, the partially observable Markov model.

2.2.1 Partially Observable Markov Model

A simple example of a discrete time Markov process is given in Figure 2.5.




Figure 2.5. A simple Markov process having three states.

At each time t , the process is in some state $x(t) \in \{x_1, x_2, \dots, x_n\}$. Given that the process is in state x_j at time t , the process will be in state x_i at time $t+dt$ with probability ϕ_{ij} where dt is fixed. Thus the matrix of probabilities Φ completely specifies the process behavior. In a real process, the state transition behavior depends on the actions taken by the operator. This dependency may be shown by $\Phi(u)$.

In general the process state is not directly observable. This can be modeled by saying given state x_j at time t the process produces the observable y_i out of a set of possible observables $y(t) \in \{y_1, y_2, \dots, y_m\}$ with probability ψ_{ij} . Thus the relationship between states and observables is specified by the matrix Ψ . This is the partially observable Markov model and may be depicted as in Figure 2.6 with circles representing states and boxes representing output tokens. There is some




Figure 2.6. A partially observable Markov process with three states and two observables.

discussion of this model in the literature [2][3][49].

We use a vector $\mathbf{f}(t)$ to represent the probabilities that the system is in each of the possible states, while $\mathbf{g}(t)$ represents the probabilities of seeing each of the possible outputs. Given $\mathbf{f}(t)$ such that $f_i(t) = P(x(t)=x_i)$ and $\mathbf{g}(t)$ where $g_i(t) = P(y(t)=y_i)$ it follows

$$\mathbf{f}(t+dt) = \Phi(u(t)) \mathbf{f}(t)$$

$$\mathbf{g}(t) = \Psi \mathbf{f}(t)$$

2.2.2 Steady State Frequencies

Of interest is the average time spent in a state under a specific control alternative. The vector of probabilities, $\mathbf{f}_0(u)$ represents the average fraction of time that will be spent in each state under action u . The solution is of interest because the expected system output under the action can be directly computed.

The solution may be obtained by solving the system of equations,

$$\mathbf{f}_0(u) = \Phi(u) \mathbf{f}_0(u)$$

By adding the additional constraint,

$$\sum_{i=0}^l f_{0i} = 1$$

the augmented system will always have a solution that is calculated according to,

$$\mathbf{f}_0 = \begin{bmatrix} 1 & 1 & \cdot & 1 & 1 \\ \phi_{11}-1 & \phi_{12} & \cdot & \phi_{1(n-1)} & \phi_{1n} \\ \phi_{21} & \phi_{22}-1 & \cdot & \phi_{2(n-1)} & \phi_{2n} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \phi_{(n-1)1} & \phi_{(n-1)2} & \cdot & \phi_{(n-1)(n-1)}-1 & \phi_{(n-1)n} \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 0 \\ \cdot \\ 0 \\ 0 \end{bmatrix}$$

2.2.3 Partitioning of a Process Model

We now show how to partition an existing model mathematically. By partitioning we mean that we will group certain states together in an effort to reduce the complexity of the model. Suppose a model having states $x(t) \in \{a_1, \dots, a_{n_a}, b_1, \dots, b_{n_b}\}$ will be partitioned so that distinctions among the b_1, \dots, b_{n_b} will be ignored. This will result in a new model having states $x'(t) \in \{a_1, \dots, a_{n_a}, b'\}$. The new state transition behavior Φ' can be easily derived and is most easily presented as a partitioned matrix,

$$\Phi' = \begin{bmatrix} \Phi_{aa} & \Phi_{ab} \mathbf{f}_r \\ \kappa_b \Phi_{ba} & \kappa_b \Phi_{bb} \mathbf{f}_r \end{bmatrix} \quad (2.1)$$

where Φ has been partitioned according to

$$\Phi = \begin{bmatrix} \Phi_{aa} & \Phi_{ab} \\ \Phi_{ba} & \Phi_{bb} \end{bmatrix} \quad (2.2)$$

and

$$\kappa_b = \begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}$$

and \mathbf{f}_r is defined by

$$f_{ri} = \frac{f_{0bi}}{n_b} \quad (2.3)$$

$$\sum_{j=0} f_{0bj}$$

As an example, we will partition a completely deterministic model having three states resulting in a model having two states and stochastic behavior. Suppose we have the system as in Figure 2.7.




Figure 2.7. Example of unpartitioned, deterministic process.

Application of (2.1) - (2.3) yields the process depicted in Figure 2.8.




Figure 2.8. Partitioned form of process in Figure 2.7.

2.2.4 Separation of Process Variables

In addition to partitioning, in which several states are grouped into a single state, one may also attempt to separate states into independently operating subsystems. This can also reduce the number of states in the model greatly, with a



Figure 2.9. Two independent subprocesses can be combined to form a single process state and behavior.

varying amount of model fidelity being sacrificed.

The inverse of the separation problem can easily be done. If two independent subsystems a and b are operating with states $x_a(t) \in \{a_1, \dots, a_{n_a}\}$ and $x_b(t) \in \{b_1, \dots, b_{n_b}\}$, and state transition behavior Φ_a and Φ_b , then the combined system state and transition matrix is given by,

$$\mathbf{x} = \begin{bmatrix} a_1 b_1 \\ \vdots \\ a_1 b_{n_b} \\ \vdots \\ a_{n_a} b_1 \\ \vdots \\ a_{n_a} b_{n_b} \end{bmatrix} \quad \Phi = \begin{bmatrix} \Phi_a \phi_{b11} & \cdots & \Phi_a \phi_{b1n_b} \\ \vdots & \ddots & \vdots \\ \Phi_a \phi_{bn_b1} & \cdots & \Phi_a \phi_{bn_bn_b} \end{bmatrix}$$

Note that while the separated system has $n_a^2 + n_b^2$ state transition probabilities, the combined system has $(n_a + n_b)^2$ state transition probabilities.

2.2.5 Relationship to Linear Systems

It is easy to show the relationship of standard linear systems models to the partially observable Markov model. However, a system with continuous variables is considered to have infinite possible states, so the state transition behavior could

never be specified in a finite state transition matrix form as we assume. Also, the representation of the system in a state transition form does not facilitate the elegant systems analysis that can be performed on linear systems. One should also remember that it is the inability of the linear model to represent many important processes that leads people to search for things like decision aids.

Linear systems theory uses the deterministic continuous system dynamic equation,

$$\dot{\mathbf{x}}(t) = \mathbf{Ax}(t) + \mathbf{Bu}(t)$$

This system may be written as,

$$\mathbf{x}(t+dt) = \mathbf{x}(t) + \left(\mathbf{Ax}(t) + \mathbf{Bu}(t) \right) dt$$

This is a deterministic model because it says that for any $\mathbf{x}(t)$ and $\mathbf{u}(t)$ there is exactly one possible value for $\mathbf{x}(t+dt)$. The same basic manipulations can be done as easily for nonlinear systems.

2.2.6 Relationship to Time Varying Systems

What is usually meant by a time varying system is that there are important system parameters that change with time or are themselves dynamic entities whose state transition behavior is not considered part of the primary system model. From our point of view, the range of possible parameter values should be added to the ordinary state vector and an augmented system should be composed on the result.

An important example of time varying systems is the failure model. In this model the process will either be in a working state or a failed state with different

state transition behavior in each case. Figure 2.10 depicts this dependence graphically.




Figure 2.10. A Markov model for a sub process that can fail.

The method of developing complete state transition behavior is similar to that of the Section 2.2.4. In doing so, a question arises as to which happens first, the failure transition or the transition within the failure? In the following, the a states represent transitions within larger states, and the b states represent these larger states. Taking the arbitrary view that the larger state transitions happen first, the process behavior can be shown to be,

$$\mathbf{x} = \begin{bmatrix} a_1 b_1 \\ \cdot \\ a_1 b_{n_b} \\ \cdot \\ a_{n_a} b_1 \\ \cdot \\ a_{n_a} b_{n_b} \end{bmatrix} \quad \Phi = \begin{bmatrix} \Phi_{a_1} \phi_{b_1 1} & \cdots & \Phi_{a_{n_b}} \phi_{b_1 n_b} \\ \cdot & \cdot & \cdot \\ \Phi_{a_1} \phi_{b_{n_b} 1} & \cdots & \Phi_{a_{n_b}} \phi_{b_{n_b} n_b} \end{bmatrix}$$

The resulting process model from such a synthesis is always considered to be time invariant, or stationary. If there are additional time varying elements in the process, then additional states need to be added to the model.

2.2.7 Relationship to Hierarchical Models

Section 2.2.4 treated combinations of independent subsystems, while section 2.2.6 presented a completely dependent state relationship. The latter typifies hierarchical structures among states in that higher level states are composed of groups of lower level states, with some transition behavior into, out of, outside of and within the high level states. Figure 2.11 schematically depicts elements arranged in a hierarchical system.

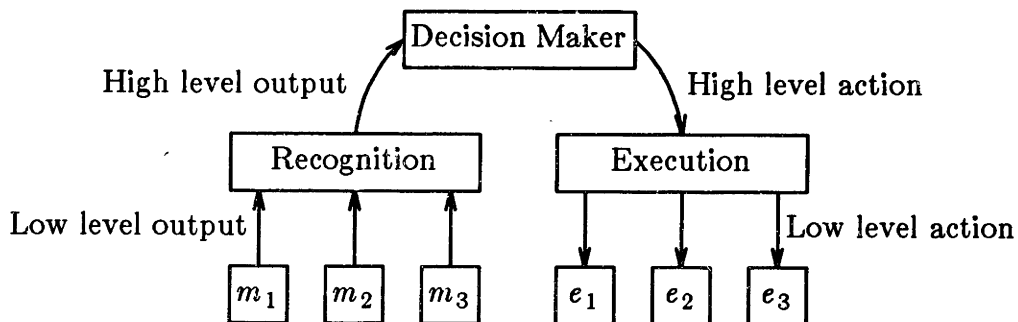


Figure 2.11. Schematic elements appearing in a hierarchical control system. The m_i are low level measurements, the e_i are low level executors.

In a system having hierarchical elements, the number and bandwidth of the interactions increases as one moves to the lower levels of the hierarchical system. The total number of elements may or may not increase; elements at the low levels may be used by more than one element in the higher levels. An advantage to using the Markov model as a basic process model is that it allows hierarchy and other realistic phenomenon to be modeled and analyzed using the same tools as nonhierarchical systems.

2.3 Normative Model of Decision Making

To judge whether the best possible decision has been made, it is necessary to have a normative description of how the decision should be made. The general model of the normative decision maker for dynamic decision making tasks is given in Figure 2.12.




Figure 2.12. Elements of the normative decision algorithm.

The remainder of this chapter will discuss this in more detail.

2.3.1 Formation of Belief State

An important part of decision making is the formation and updating of a belief state. In the normative decision maker, this will consist of a probability distribution over all possible states of the process. We represent this by $f^-(t)$ or $f^+(t)$ as in figure 2.12, or simply $f(t)$. This belief state is sufficient to represent the knowledge of all past observations about the state of the process [3].

If the process can be in only a finite number of states, the belief state will be a vector of probabilities. If the process has continuous state variables and therefore infinite possible states, the belief state is a likelihood function defined on the state space. Normally, the value of the belief state will change as time progresses and additional observations are made. Before any measurements are made, the

steady state frequency vector \mathbf{f}_0 or the uniform distribution over possible states may be used as *a priori* belief states. The former is preferable because it takes additional information about system parameters into account.

2.3.2 Updating Belief State in Absence of Observation

If the observables are ignored for a period, one's beliefs about the process state should still change consistent with his model of the process dynamics. This is done simply for the finite state case. The expected state before observation $\mathbf{f}^-(t)$ is computed from the last updated estimate $\mathbf{f}^+(t-dt)$ and the process dynamics Φ according to

$$\mathbf{f}^-(t) = \Phi(u(t-dt)) \mathbf{f}^+(t-dt)$$

This takes into account the process dynamics associated with the particular action taken at time t .

2.3.3 Updating Belief State Through Observation

In addition to the ongoing evolution of the process state, the observation of additional process outputs should also influence the belief state. This is done according to a normal Bayesian update and assuming y_i was observed at time t may be written as,

$$f_j^+(t) = \frac{\psi_{ij} f_j^-(t)}{\sum_k \psi_{ik} f_k^-(t)}$$

Note that if there is uncertainty in the observation so that the observation is given as a distribution $\mathbf{g}(t)$ so that $g_i(t) = P(y(t)=y_i)$, the update may still be computed as,

$$\mathbf{f}^+(t) = \frac{\text{diag}(\mathbf{g}^T(t)\Psi) \mathbf{f}^-(t)}{\mathbf{g}^T(t) \Psi \mathbf{f}^-(t)}$$

2.3.4 Weighing Alternatives

Given a belief state, one must weigh the various actions available and select one as a decision. The basic normative method for selecting an action is maximizing the expected incremental value over the possible actions [15]. To do this we add a constant reward function \mathbf{r} . The element r_i of reward function \mathbf{r} represents the payoff when the system is in state x_i .

If one only considers the effects of the present action on the next state of the system, the computation of the optimal action u_o is straightforward and is given as,

$$u_o = \underset{u_i}{\text{argmax}} \mathbf{r}^T \Phi(u_i) \mathbf{f}(t) \quad (2.4)$$

This will be the exact solution when the decision problem will end on the following state. However, for an ongoing decision problem, one must consider the effects of the present action on all future time and select an action accordingly. To do this, a selection policy must be specified so that future actions, future states, and future rewards may be predicted.

2.3.5 Approximation to Normative Selection Rule

As we have set up the problem, rewards are received for being in a specific state for a single timestep. To approximate the normative decision rule future rewards must be accounted for. If one delays filing one's tax return on February 1 as compared with April 10, the immediate reward has the same value which is one day without doing taxes. The effective reward differs, in that as the tax

deadline approaches, it is more likely that a day's delay will result in paying tax penalties for late filing.

It is assumed the resulting structure of the effective reward consists of additive linear functions and can then be used as a basis for making a single-state prediction comparison as in (2.4).

Thus the effective reward \mathbf{r}' consists of the original reward \mathbf{r} plus the future rewards expected given the present state. By setting one equation for each state, \mathbf{r}' may be formed by solving the system of equations,

$$r'_i = r_i + \mathbf{r}'^T \Phi(u'_i) \quad (2.5)$$

where u'_i is the optimal action given the process state is x_i .

The steady state frequencies obtained by using control law (2.4) with (2.5) can be approximated by solving the approximately equivalent dynamic process under an *ad hoc* stochastic controller

$$P(u_i) = \mathbf{r}'^T \Phi(u_i) \mathbf{f}_0 / \sum_{j=1}^m \mathbf{r}'^T \Phi(u_j) \mathbf{f}_0 \quad (2.6)$$

and setting the equivalent process dynamics to be

$$\Phi_{equiv} = \sum_{i=1}^m \Phi(u_i) P(u_i)$$

In the finite state case, there are a finite number of possible action functions \mathbf{u}' . Each action function should be assumed and the system (2.5) solved to determine the effective reward function under that action function. Then the approximate steady state frequencies under this control policy can be computed and the expected average reward for that action function can be computed according to (2.6).

Once this is done, the action function yielding the highest expected average reward can be chosen, and its corresponding effective reward function determined. The most desirable action can then be chosen according to (2.4) using \mathbf{r}_e in place of \mathbf{r} .

2.3.6 Uncertainty in Belief State

As a measure of the uncertainty in a belief state, we will use the common information theoretic measure of uncertainty in a distribution [53]. For the finite state case the uncertainty I associated with probability distribution \mathbf{f} is,

$$I(\mathbf{f}) = \sum_{i=1}^n \frac{1}{f_i} \log_2(f_i)$$

The uncertainty in the *a priori* belief state as represented by the steady state frequencies is in some way a measure of the complexity of the process. This is a basic measure of the minimum average number of bits of information necessary to encode state transition behavior over long periods of time. It will always be less than or equal to the number of states of the process.

2.3.7 Uncertainty Growth in Absence of Observation

In the absence of observation and over time one's uncertainty about the state of a process will grow. More precisely, since the belief state will tend to the steady state frequency distribution \mathbf{f}_0 , the uncertainty in the belief state will tend toward the uncertainty associated with this state $I_0=I(\mathbf{f}_0)$. As a measure of the rate at which uncertainty will naturally grow about a process state over time and hence an important characteristic value of the dynamic process, we define the uncertainty growth rate I_p to be,

$$I_p = \sum_{i=1}^n I(\phi_i) f_{0i}$$

This measure is derived by weighing the expected uncertainty just after a state $I(\phi_i)$ is known with the *a priori* probability of seeing that state f_{0i} .

2.3.8 Uncertainty Reduction Through Observation

When an observation is made, the uncertainty about the process state will normally be reduced. In analogy to the previous section we define another important characteristic value of a process, the average uncertainty reduction rate I_q . This is defined as

$$\begin{aligned} I_q &= E(I(\mathbf{f}^+(t) \mid \mathbf{f}^-(t)=\mathbf{f}_0)) \\ &= \sum_{i=1}^m I \left(\frac{\text{diag}(\bar{\psi}_i) \mathbf{f}_0}{\bar{\psi}_i \mathbf{f}_0} \right) \bar{\psi}_i \mathbf{f}_0 \end{aligned}$$

This is a measure of the rate at which ones knowledge can be expected to increase about a process state, given that you know nothing about it at present. It is defined in this way to characterize the stochastic observability of a process.

2.3.9 Finiteness of Symbols and Reduction of Belief State

Complexity

For the finite state tasks, the normative belief state is a vector of probabilities, one for each state. This can usually be implemented on the digital computer. When, as in most important cases, the system has continuous variables, the normative belief state consists of a likelihood function that is even more difficult to implement.

We may wish to study how belief states change for normative or real decision makers, to determine theoretical performance levels. This involves forming a belief state over the decision maker's belief state, and so is unwieldy even for the case of the simplest finite state systems. What must be done in these cases is to use approximations to the belief states that as much as possible retain the behavior of the exact belief state. A typical approximation in the continuous case is to retain only the mean and variance of the belief state thereby substituting two values for the infinitely valued exact belief state. In the same way, for a finite state system, one may develop a decision making model that is itself a finite state automaton. In such a case the steady state frequencies of the combined process plus decision maker system should be computable.

The problems of human decision makers in this regard are obvious. The human has limited short term memory resources that he can devote to maintaining a belief state consisting of a distribution over the process states, if he is to be a normative decision maker. It is therefore not surprising that he may simplify his picture of the world in his attempts to control it. His long term memory appears to be limitless, but it takes time to acquire new knowledge and symbols. Within these limitations, it is possible to behave normatively by having a different symbol in long term memory for each possible belief state. The successor of each belief state for each possible observation must also be stored. The accumulation of a large body of context specific information such as this is what long training periods and years of experience accomplish. To be a normative decision maker with the normal human processing limitations requires either a differentiated sym-

bol in long term memory for each possible belief state or extremely long periods between decisions and lots of pencils and paper.

2.4 State Estimates as a Decision Aid

For continuous processes, it is not practical to compute a complete estimate of the process state or communicate such an estimate to the human operator. However, an approximation to the distribution may be represented by a moment expansion. Such an approximation can serve as the basis of a decision aiding system by augmenting the human's own abilities to estimate the state.

For example, if the state of a process is represented by a vector having n components, then a first order approximation to the belief state will be a point estimate vector having n elements. If the moment expansion is used, the value of the approximate belief state will then represent the expected value of the state vector. This is the approximation typically used in control systems. A second order approximation consists of a vector as a point estimate and a $n \times n$ covariance matrix to represent the uncertainty of the estimate. In estimation theory, this second order expansion is often used.

By using an expansion of the beliefs higher than first order, we hope to improve decision making performance. Since this is a significant increase in information for most systems, it remains an open question how such information is to be used in any practical way by the human operator, or if it can be used at all. Later, we show that uncertainty is useful for simple systems, but is seldom used by human operators for more complex systems.

3. Considerations in the Design of Decision Aids

We now present a qualitative analysis of the decision aid design problem. For the moment we assume that such an analysis is useful in and of itself, even though it cannot be backed up with corresponding quantitative statements and analysis.

3.1 Statement of the Decision Aid Design Problem

The decision aid design problem, like any design problem, consists of specifying a set of characteristics for a decision aid that ensures that decision making performance will be improved. To make his selections, the designer may do an analysis of the man machine system for which it is intended. We will try to explain the scope and range of possible analyses in this chapter.

3.1.1 Element Characteristics and Performance Measures

We begin with the question, "What are the interesting characteristics of an aided man machine system and how do they interact to produce better or worse system performance?" To begin with, there are three distinct components of the aided system. These are the controlled process, the human operator, and the decision aid. Each will have a set of characteristics that will influence the quality of the decision making made by the human operator. The goal of this chapter is to identify an important set of characteristics for analysis and their interactions.

Given a specific combination of man, machine, and aid characteristics, what is the resulting system quality? Ideally, system quality is measured by taking a long term average of the system output when the combination is used. To understand interactions among characteristics means to be able to predict the system quality of combinations without actually constructing the system.

3.1.2 Fixed and Selectable Characteristics

Throughout this chapter we will discuss characteristics. The characteristics of the decision aid may in general be selected by the designer. The human and task characteristics are generally fixed before the design problem is considered and it is the role of analysis to bring all the characteristics to the consciousness of the designer. The designer chooses decision aid characteristics out of those that are possible to maximize the average system output. What distinguishes one decision aid design problem from another is the unique set of fixed characteristics present in the controlled process.

3.2 Process and Task Characteristics

Though we have slightly different conceptions of what the controlled process is and what the decision maker's task is, we treat them here together as either the process or the task. Generally speaking, process refers to the process state transition and output matrices, while task refers to the utility function over process states. We use process and task almost interchangeably, and by either we mean anything that is outside both the human operator and the decision aid.

3.2.1 Interdecision Rate

As a predictor of the success and character of a man machine system, we consider the basic time available to make decisions to be an important variable. All other things being equal, longer decision times will generally result in better decisions and higher system output because the human will be able to apply something closer to the normative decision procedure on the problem. Exceptions will occur if the operator becomes bored between decisions, or if he is controlling a very slow system.

Examples of processes where decision times are short are landing an airplane and crossing a street. Slow systems include steering a ship and controlling the money supply. We may use the fastest eigenvalues of a system as a rough indicator of how often decisions will have to be made when controlling it. In a typical supervisory control system, the inner "loops" have much shorter decision times and are automatically controlled while the slower outer loops are controlled by the human.

3.2.2 Known and Unknown Dynamics

Sometimes the process dynamics are precisely known while in other cases they have not been modeled very well. Unmodeled process dynamics have already been cited as a major reason for having the man in the loop at all. Important examples of poorly modeled systems are economic systems, political systems, and military systems. Since these systems are composed of competing elements, we can expect them to remain unpredictable and incapable of ever being well modeled. This is because if it is possible to model a competitor, it should then be possible to exploit him. It follows that as long as there are competitors, tomorrow's stock market price or the next move of a capable adversary will be for all practical purposes unpredictable.

Systems that appear to be well modeled, such as power plants, may also be considered poorly modeled if the entire model was not used in generating the control policy. The occurrence of failures is generally considered a change in the controlled process - in essence the failure is not a part of the model used to determine moment to moment control actions. In these cases the human can presumably make better decisions than the automatic equipment that is suddenly making decisions based on a faulty model.

In a system that is accurately modeled, we may wish to remove the human from ongoing decision making as much as possible since the human is unreliable as a simple information processor. However, in a system for which the model is not known explicitly, we may be willing to trade this slight unreliability for the robustness that a human can offer in new and unfamiliar situations.

3.2.3 Simple and Complex Systems

We all recognize complexity in tasks and processes. Controlling a power plant is more complex than riding a motorcycle. New cars are more complex than older cars. However, there is little agreement on exactly what constitutes a complex process.

Factors affecting complexity include the number of possible states, actions, and outputs of a system, the degree of interconnections among subsystems, and the degree to which effects of actions are not obvious. A decision aid may help an operator deal with complexity by reducing apparent complexity. Removing excessive degrees of freedom by controlling inner loops, helping the operator predict the effects of his actions, and coordinating control to decouple subsystems are ways in which decision aids might be helpful.

A possible characterization of complexity is the information in the steady state visitation frequency vector $I(\mathbf{f}_0)$. This measures the size of the comprehension system necessary to model the system before observations have been made when uncertainty about the process state is at a maximum. A system composed of two completely independent subsystems having complexities I_a and I_b will result in a system having complexity $I_{ab}=I_a+I_b$ while one composed of two dependent subsystems will have at most $\max(I_a, I_b)+1$ bits of complexity.

3.2.4 Predictability of Process State

Various processes, though they may be well modeled, are predictable to varying degrees. The motion of Halley's Comet is well modeled and extremely

predictable. Weather predictions are inaccurate after a few days despite a thorough understanding of the principles involved in weather formation. The stochastic elements or uncertainty in the process model produce unpredictability, and there is no way to remove all uncertainty in a model.

The process uncertainty I_p reflects the uncertainty inherent in a process model. Clearly a decision aid that conveys the uncertainty of a prediction along with the prediction is preferable to one that gives the illusion of predicting the future with perfect accuracy. There are other issues of whether the operator will accept such information or even whether he has time to digest it that may be preventing the widespread use of probabilistic displays, however.

3.2.5 Observability of Process State

Inability to observe or measure process states may also confound the decision maker. For example, the law of supply and demand is a basic principle for many economists. Yet one can scarcely measure supply, let alone demand, and if entirely new markets are discovered or created, the modeling problems are even worse. Price is cited as a visible measure of the balance point between current supply and current demand, yet in the stock market the price of a share is hardly a constant, steady quantity.

Sometimes a computer may alleviate the human of tedious calculations that he would otherwise do to estimate variables that are hard to observe. By relegating these tasks to the computer, one is assured that the computations are carried out precisely and repeatedly. Mendel [38], for example, is developing a computa-

tional method for combining information from various sources. In the absence of an automated system to do this, an operator would certainly combine the information anyway, though there is no assurance that he would do so in an appropriate way. By using a computer to do these tasks, one has control over how this combination takes place.

3.2.6 Single Time, Known Final Time, and Ongoing Tasks

Tasks may be characterized by how many times the decision must be made. There are three important cases: single time, known final time, and ongoing decision problems. Theoretically, this determines how far into the future ones expected value calculations must be computed. Section 2.3.4 presents normative selection rules for the single decision problem; Section 2.3.5 modifies this for the ongoing decision problem.

An example of a single time decision is buying stock, when the decision is, "When should I buy stock?" Once stock is purchased the decision is not made over again, though at times an investor may wish it could be. A finite time decision problem is competing in an election. In this problem the actions are dependent on the exact election date as well as the perception of the opponents strategy and relative standing. Infinite time or ongoing decision problems are called regulation problems. A typical example is the control of a power plant where the operator must continually reassess the state and adjust his inputs.

3.2.7 Forcing a System and Capturing a System

There is a qualitative difference between systems in which the state is forced

through various actions and a system where the process state runs freely, and then is captured when it enters a desired state. Largely this is a perceptual difference in the way we think about problems. Experiments one and three reported in Chapter 4 are forcing problems, while experiment two is a capture task. One can also see this schism as a regulate versus track dichotomy. In regulation, inputs are continually being adjusted to attempt to bring about a new state, while tracking involves a lot of watching and a few inputs when the time is right.

Examples of forced systems are common. Driving a car, determining exchange and interest rates, and influencing peers through manipulative social behavior are some. On the other hand, businessmen making investment decisions, students working long hours when they happen to want to study, or politicians who seize a naturally occurring shift in public opinion to advance their own causes are examples of operators engaged in capturing behavior.

3.2.8 Target State and Avoidance States

The structure of the utility function also determines how an operator approaches a decision problem and whether his natural ways of coping with other issues such as complexity will likely lead to good or poor decision making. At one end of this range are target state tasks in which a few of the states are rewarded while the rest are unrewarded. At the other end are tasks dominated by avoidance states where some of the states are penalized and the rest are uniformly unrewarded. Examples of utility functions producing these tasks are given in Figure 3.1 .




Figure 3.1. The shape of the utility function determines if the system has target states or avoidance states.

An example of a target state task would be traveling to a specific location. The decision maker is considered unsuccessful in any of the myriad ways that lead him to any but the desired location. An example of a task having avoidance states is keeping a troubled business afloat. In the mind of the maverick business executive, success is equally sweet for any state in which he manages to keep things going, and his decisions will be commensurate with this attitude.

3.2.9 Forgiving and Unforgiving Reward Structure

In addition to the shape of the utility function, we consider the peakiness or severity of the extended regions to be an important characteristic of the task. Figure 3.2 shows two different utility functions of different severity. The function on the left has avoidance states, but the right hand function *really* has avoidance states. Tightrope walking in the circus is characterized by avoidance states, but tightrope walking without a safety net is also characterized by extreme severity.

Because the utility function is only unique to within a positive affine transformation, it is difficult to assign a precise meaning to the concept of severity as we have presented it. Approximately what we mean is that most states are




Figure 3.2. Two utility functions of varying severity.

undifferentiated with respect to utility, and a few have very large utilities compared to this nominal differentiation. Professional golf is an activity with target states and a high degree of severity in a positive sense. Virtually all golfers make no money or pocket change, while a handful earn much more than a living from it. It is not the large salaries *per se*, but their size in comparison to all the other golfers that makes the problem severe.

Severity is difficult to cope with from a design standpoint because it brings about evaluation problems. If we are trying to evaluate a golf training program for example by using average golfing salary as a measure of training program performance, the presence or absence of a wage earning caliber golfer will completely overwhelm the rest of the data. In evaluating decision aids, measures must be chosen that do not have this severity problem.

3.3 The Fixed and Partially Known Human Operator Characteristics

We now review what is known about the fixed characteristics of the human operator as a decision maker.

3.3.1 Simplification Models

Chapter 4 will present some simplification models of the human operator. These are based on experimental evidence and the idea that the human decision maker imperfectly performs the normative decision making algorithm. Any element of the normative model can exist in the human in a subnormative form. The basic elements and the likely ways in which they will be subnormative are,

- Task model. The human decision maker may have misconceptions about the system behavior that are convenient to hold or familiar from other experience.
- Utility function. The human decision maker may view the problem as a target state problem when it is not relative to the utility function of society or the management of the company. This simplifies his action selection computations.
- Belief state. The belief state of the decision maker will be highly simplified compared with the normative belief state. Most likely it will be a point estimate and no more.
- Prediction process. This may be inaccurate or non-existent, especially if decisions are thought to occur at a much higher rate than the basic process rate.
- Updating process. Observations will not be combined optimally with existing beliefs. Systematic biases will enter these computations .
- Selection process. Not enough attention will be paid to the future consequences of the present action. The operator will apply a single state

lookahead rule. For example, a restaurant may serve small portions for large prices because they fail to realize the importance that present customer satisfaction has on their future profits.

- Execution. Occasional errors will be made executing intended actions. For example, erroneous data will occasionally be entered.
- Perception. Occasional errors will be made in perception. Repeated errors perceiving the same observable may be made if expectations about what will be observed are high.

3.3.2 Models from the Literature

There are many models of human decision making in the literature. Most of the important ones are referenced in Section 1.3.1.

Rasmussen's skill, rule, and knowledge based hierarchy^[45] categorizes operator behavior into these three distinct classes. While skill and rule based behaviors are the dominant behaviors exhibited by inexperienced operators, knowledge based behavior occurs in the expert operator and involves the accessing of facts about the system to predict the results of actions under consideration, possibly for circumstances that have not occurred before.

Workload models of the human operator have proved to be important in aircraft pilot studies. The workload model derives itself from the concept of the human as a limited capacity information processor, or the computer analogy.

3.3.3 Training of Human Operators

With sufficient training, human beings can perform remarkable tasks. Therefore, the human's abilities should not be considered absolutely fixed by the decision aid designer. The human, through training, can accustom himself to what seems like any display or condition. This is perhaps one of the confounding factors in the study of man machine systems. Whatever a subject is given in the way of an interface, he is likely to be able to use it proficiently after enough training. Of course if there is a choice between two designs that result in equivalent decision performance, the choice requiring less specialized training would be preferred.

3.4 Decision Aid Characteristics

We now look at several ways in which decision aids may be characterized and some guidelines for choosing desirable characteristics.

3.4.1 Integrated and Separate Displays

An integrated display is one in which several variables can be displayed as elements of a larger picture or in the same physical space. What we call a separated display is one where each value receives its own demarcated space for display. The artificial horizon gauge in a cockpit is an example of an integrated display. The normal dashboard of a car consisting of several independent gauges is a separated display.

Within integrated displays, we distinguish between pictographic and nonpictographic integrated displays. In the former, the graphic image is made to look like the values that it depicts, while in the latter, no effort is made to make correspondence between image and reality.

Integrated displays have the advantage of conveying a larger amount of information per unit area of display space, and presumably per unit of the operators time. This relies on the assumption that the operator can process an integrated display in parallel. The work on Chernov faces in which up to 21 variables can be displayed simultaneously would suggest that people can process integrated displays in parallel given training. Hence integrated displays may be most useful when decision times are short, or for repetitive tasks with high information requirements that allow for special training of the operators.

One disadvantage of an integrated display is that if a single variable measurement is desired, it must be extracted from the integrated display and this takes mental effort or special training. For a pictographic integrated display, it may also be hard to make accurate readings. Thus for a target state task with longer decision times, a separated display may be preferable.

3.4.2 Forward and Backward Looking Displays

In the forward looking display, predictions about the process state are generated assuming some control law and then displayed. This is the normal predictor display [34]. The predictor display has been shown to be useful for extremely slow systems like ships, or slow process control situations.

The backward looking display will show process states back from some reference point. We distinguish two types of backward looking displays. In the first, a trajectory of states is displayed that lead to a desired state or target state, and so may be most useful in tasks having target states. In an airplane landing task that is a target state task, a landing envelope could be displayed to show desirable trajectories leading to successful landings. Note that this puts the operator in a rote procedure mode because he must only follow the path to succeed.

In the second type of backward looking display, the history of the process state knowledge is traced out backwards from the present time. This allows the operator to do an *ad hoc* statistical summary of the data and predict the future in a statistical sense. This may be useful for unpredictable problems. The display

of stock market prices over the last year, the history of rainfall, or the progression of a forest fire are backward looking historical display conveying little more than statistical information about the variable of interest.

3.4.3 Probabilistic and Mean Value Displays

To behave normatively, one must maintain a belief state that is a probability distribution over all possible states of the process. But in practice, seldom is more displayed than the mean value of an estimated variable. Even if the uncertainty information is displayed, it is probably not used by the human decision maker. Hence the decision aid designer has the choice of displaying additional probability information to the operator or displaying only mean value information.

If probabilistic displays are used, they will certainly add to the operator's workload when compared with a mean value display. But if the information is integrated, the human may be able to process it just as quickly by processing it in parallel. Yet operators are not practiced in using this type of information so training would be necessary to expect improved decision making.

Probabilistic information is important in unforgiving systems where throwing away small probability masses can be dangerous because of the large utilities involved. They may also be useful in systems where observability or predictability are low, especially if predictor displays will be used and time constraints are low.

3.4.4 Forced and Unforced Information

A designer faces the choice of forcing information on the operator or merely having it available for the operator when he would like it. Information residing in these two forms we will call forced and unforced information.

For forced information, there is no time spent bringing up the information, so this is more suitable for time critical applications. Second, there is a smaller possibility that the operator will forget the information is available, so important information should be presented in a forced format. As a disadvantage, there may be the problem of filtering also present for integrated displays. If information that is not desired is also presented, the operator will filter it out to extract desired information.

When a system contains unforced information, there is a higher total capacity for information that would overwhelm the operator if it were all forced. A disadvantage is that it takes longer to find a specific piece of information. In addition, the operator may forget completely that needed information is available at all, especially if access to the information is infrequent. How many of us remember the city population chart in the dictionary? An unforced format is not suitable for those variables that are critical to successful task decision making.

3.4.5 System Flexibility

Similar to the forced/unforced issue is the flexibility issue. A decision aiding system may be designed to be flexible to varying degrees. A flexible system may allow entry of new data, development of new procedures, or even a restruc-

turing of what information is forced.

Certainly the place for flexibility is in a system that is poorly modeled. Flexibility there would allow the operator to develop and enter subprocess models as he himself discovers them. However systems that are flexible are generally time inefficient, difficult to learn because of the presence of meta-instructions, and prone to all sorts of unexpected behavior due to person *A* changing person *B*'s configuration and other odd problems. "Being given enough rope to hang yourself" is a common expression referring to what happens when there is too much flexibility in a system. Inflexible systems, being much more time efficient and less prone to problems, are much preferred in unforgiving situations.

3.4.6 Complete and Incomplete Knowledge

Of course no model of a real process is ever completely accurate, but a decision aid may or may not contain all of the available information about a controlled process. All systems will benefit if the decision aid is operating on information that is complete in this sense. However, an important situation that can benefit from incomplete systems is an overly complex system with short decision times. The time constraints may prevent using a complete model, but if the decision aid controls inner loops where it needs to only know a few things, the effective complexity that the human sees can be reduced.

For a complex, poorly modeled system with low time constraints, incomplete models and flexibility can be a valuable combination. An example is economic systems that are slow, complex, poorly modeled systems. Despite the

incompleteness of the models available, the Fed continues to adjust interest rates frequently and fearlessly and is flexible enough to base its decisions on whatever theories are currently fashionable.

3.4.7 Hierarchical and Sequential Concepts in Decision Aids

The concept of hierarchy was introduced in 2.2.7. Presumably, hierarchical elements will be of use in the decision aid to the extent that the system is itself hierarchical. An important part of the archetypical supervisory control system is the hierarchical element that succeeds in implementing the human's control behavior in automatic equipment, thereby reducing complexity from the human's point of view.

There is no place for hierarchical control in the control of very simple processes, and the need is reduced for single time decision tasks. There is a semantic difficulty in saying this, however, since the act of implementing a control law such as $\mathbf{u}(t)=\mathbf{K}\mathbf{y}(t)$ is a hierarchical action, despite the apparent simplicity of the stated action.

3.4.8 Knowledge Driven and Procedure Driven Systems

The difference between knowledge based systems and procedurally driven systems is a dichotomy that is reversed in the power industry, but for which little substantial basis exists. A decision maker who operates on knowledge deduces procedures as an intermediate step to action, while the procedurally driven decision maker immediately, and blindly, invokes stored procedures. The knowledge driven decision maker will be able to select reasonable actions in unfamiliar cir-

cumstances, while the procedurally motivated decision maker will not. The same basic properties will be present in decision aids based on knowledge versus procedures. A knowledge based system will certainly be most useful in complex control problems, though the construction of useful knowledge based systems for the control of processes is an open research question.

3.5 Systems Considerations in Decision Aid Design

3.5.1 Human and Computer in Single Time Decision Task

Figure 3.1 depicts the role of human and computer in a single time decision task. On the left is a flow diagram of information and processes that would be used by the human in the decision task, and to the right is the corresponding diagram for a computer doing the task alone. There is a basic correspondence between the human elements and the computer elements.

Generally we feel that the computer is superior to the human in the sense, estimate, and act processes, while the ability to form plans automatically is not understood or perhaps will not be trusted to a computer. When a computer aid is available, the computer may be assigned the tasks at which it excels, while the human does only planning, which is considered the deep decision making.

What we must remember and an important part of what this diagram shows is that although both human and computer have an estimate or a plan, the representation system used by each is completely different. This means that if the task is to be shared and therefore information is to flow between human and computer, interfaces will be required to allow this communication. The two places

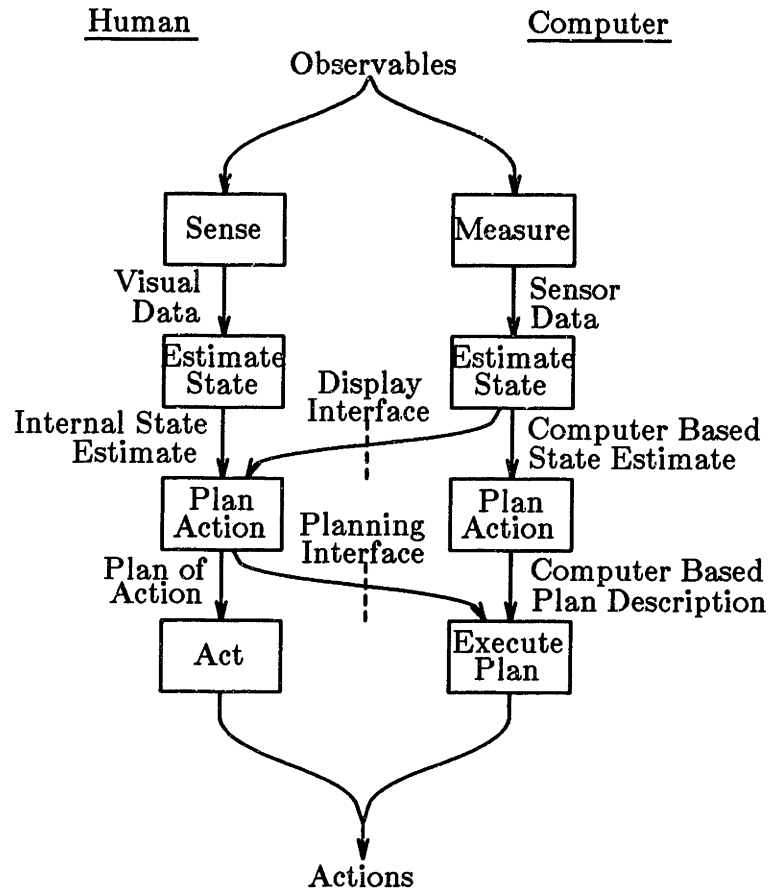


Figure 3.3. Role of human and computer in single time decision task.

where we would likely find interfaces are between computer and human estimates, and between human and computer plans. These interfaces will represent additional workload, complexity, and potential confusion and error. In systems that are already complex or have limited decision time, these additional effects of the decision aid must be kept in mind and weighed against their other theoretical benefits.

3.5.2 Perception and Action as Mediators Between Forms of Information

3.5.3 Assessing the Value of a Decision Aid

As we stated in Section 1.1.4, there is an inherent problem in assessing the value of a decision aid. Roughly, if a normative procedure exists so that decisions can be judged independently, the normative procedure should be used in place of the human. In cases where a normative procedure cannot be derived, the value of a decision cannot be individually assessed.

The value of a decision is the expected reward derived specifically from making that decision. Ultimately, the long term average system output is the best measure of the value of a decision aid. This can be approximately measured in an experiment by measuring the system output directly and averaging over the experimental trial. An action will have an effect on a system that will last a shorter or longer time, with the effect being longer and more multifaceted for more complex systems. For an experimental trial that is much longer than the typical time for which actions have an effect, the sample average will be a close measure of decision making performance *for that set of observables*. The sample average will have a high variance for systems that are unforgiving, however, so it may be of little use in these cases.

Suppose we are conducting an experiment consisting of a one time capturing task, and the subject has made his decision. We would like to evaluate his decision. Certainly we will judge the decision in light of the information that he had at the moment of the decision, so we must take into account the process state knowledge at the decision time and whatever this knowledge meant about future process states. Figure 3.2 summarizes the process state information that might be

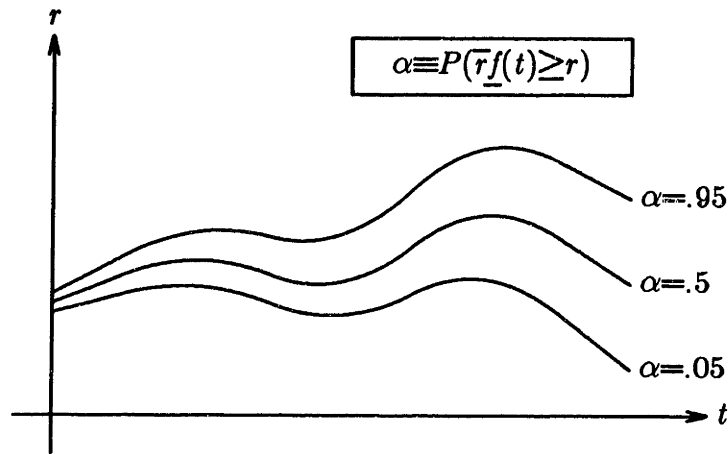


Figure 3.4. A summary of process state knowledge combined with the utility function extending into the future. Bounds represent 5 to 95 percent certainty region.

available at the moment of the decision.

A one time capturing decision will be optimal if the expected value of the decision exceeds the expected value of delaying the decision for all points into the future. Therefore, the expected loss in making the decision now is the difference between the expected value of making the decision now the best possible expected value in the future.

Expected value is not the only criterion that could be used, however. One may wish to be 95 percent certain, for example, that no better opportunity will arise in the future. By selecting different levels of certainty in Figure 3.4, one can choose a more or less conservative selection rule. The value of the decision can be parametrically determined with respect to the cutoff percentage.

4. Three Experiments in Human Decision Making

We now report on three experiments conducted to explore human decision making in dynamic decision making tasks. The experiments were designed to study the uses of normatively derived state information in decision aids as well as gaining general insight into human decision making.

4.1 Introduction to Experiments

Three experiments were carried out that are reported in this chapter. We first give the experimental description common to all three experiments. Later sections will treat the particular aspects of each experiment separately.

4.1.1 Forced Pace Dynamic Decision Making Task

In each experiment, subjects were asked to control a dynamic process by selecting from possible actions on a repeated, ongoing basis. The task was a

forced paced, discrete time task with the time between decisions ranging from 0.5 seconds to 2.0 seconds. Always, subjects were given an accurate description of the process dynamic characteristics and an explanation of the scores received for being in each of the states.

Subjects were seated in front of computer generated displays containing various state information. At the end of each decision interval, the subject indicated his decision with of a switch, and information about the new process state was immediately displayed. For a given display configuration or experimental condition, the subject would make between 50 and 200 successive decisions, and only then would his score for that trial be displayed. Training consisted of a sample experimental run in each new condition until the subject felt comfortable with the task before data was recorded. Because of the simple nature of the process dynamics, this would typically require a single training run before subjects reached a steady level of performance.

4.1.2 Simulated Process Having Few States

The partially observable Markov model of Section 2.2 was used as a process model, and it was simulated on the digital computer. This allows the experimenter to know the true behavior of the process. Processes with few states were used so an approximation to the normative control algorithm could be calculated. This normative control procedure allows the expected value for each decision to be calculated. The computer recorded each decision made for later analysis.

4.1.3 Display of State Estimates as a Decision Aid

We consider there to be three primary sources of information about the process state. First, there is direct knowledge of the state. Though this does not occur in the real world, it is useful in experiments as a control condition, or it may help the subject learn the processes dynamic behavior. Second, there are the process observables. These are the only source of information available normally, and both human and decision aid have only this information to go by in their decision making activities. Finally, we consider a decision aid based on providing a normatively derived estimate of the process state to be a third source of information. Though the decision aid really has no more information about the state of the process, it does represent a convenient summary of the information from past observables combined with the process and output models. All three forms of state information were used in various combinations throughout the experiments.

4.2 Experiment One: Three State Tracking Task

4.2.1 Qualitative Process Characteristics

A task was set up having three states, three outputs, and two possible inputs. Under input u_1 , the state would tend to move in one direction from x_1 to x_2 to x_3 . Under input u_2 , the state would move in the other direction from x_3 to x_2 to x_1 . The outputs suggested which state the process was in, but were sometimes wrong. For each time the state was in x_2 , the subject received one point while states x_1 and x_3 were unrewarded. Under some conditions normatively derived state information was also displayed.

The cover story given to subjects was as follows. "A certain frog can be in one of three buckets [the process states]. A carrot can be positioned on one side or the other [two actions] to coax the frog in one direction or the other. Under most conditions the frog cannot be seen, so a dog is used that points at a bucket but is occasionally wrong [positively correlated observables]. Your goal is to keep the frog in the center bucket [reward function]. Occasionally a nerd is available to calculate the present probability of the frog's location [decision aid], and he displays these as a bar graph."

4.2.2 Quantitative Process Characteristics

The process characteristics consist of the two Φ matrices, the output matrix Ψ , and the reward function, \mathbf{r} . Because this was an exploratory study, several different matrices were used parameterized by two variables p and q . The characteristics used were,

$$\Phi(u_1) = \begin{bmatrix} 1-p & 0 & 0 \\ p & 1-p & 0 \\ 0 & p & 1 \end{bmatrix} \quad \Phi(u_2) = \begin{bmatrix} 1 & p & 0 \\ 0 & 1-p & 0 \\ 0 & 0 & 1-p \end{bmatrix}$$

$$\Psi = \begin{bmatrix} 1-q & q/2 & q/2 \\ q/2 & 1-q & q/2 \\ q/2 & q/2 & 1-q \end{bmatrix}$$

$$\mathbf{r}^T = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

4.2.3 Experimental Design

Three subjects who were graduate students in engineering participated in the experiment. They were given the cover story in Section 3.2.1 as well as the

numerical process characteristics. A schematic form of the display showing the possible outputs, the actual output, and the decision aid is given in Figure 4.1.

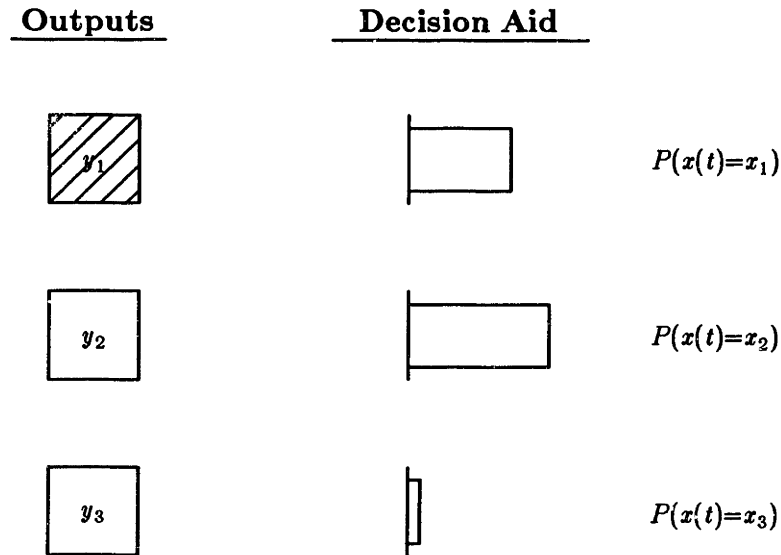


Figure 4.1. Schematic display showing an observed output and the decision aid.

The time between decisions τ was varied and two display configurations were used, one showing only process observables and one including the decision aid. Section 4.2.2 gave the process transition and output matrices parametrically. The range of values for τ , p and q used and the resulting uncertainty growth and reduction rates I_p and I_q are as follows:

p	I_p	q	I_q	τ (sec)
0.055	0.187	0.026	0.131	0.267
0.160	0.333	0.106	0.391	0.533
0.371	0.418	0.224	0.651	1.067
		0.410	0.911	2.011

During training subjects practiced the two experimental conditions. They also practiced with a visible state condition and a state plus observables condition.

This allowed them to develop a sense of the process dynamics. Protocols of the subjects were recorded and reviewed later.

4.2.4 Usefulness of Decision Aid

To determine whether the decision aid was of any use to the subjects, we scored the decisions made by the subjects under each condition. A simple scoring rule is to take the score received in the experiment as displayed to the subject at the end of each trial. A better rule is to calculate the expected value of each decision and sum these results. For infinitely long experiments the two methods will yield the same results relative to each other. For short experimental trials, the latter method will produce data with a much lower variance.

Figure 4.2 shows the results of comparing the aided to unaided cases using the two different scoring functions. As expected, the scores were improved in the aided case over the unaided case. This is more clearly shown with the expected loss as a scoring function that has a greater ability to differentiate the scores for short experimental trials.

4.2.5 Effects of Process Parameters on Decision Making

Figure 4.3 give the results of an analysis of variance on the experimental data and summarize the results in a tabular format. The raw trial scores have been used in this compilation. The table shows that the presence of the decision aid and the process uncertainty have significant effects on the decision making performance. The uncertainty of the process also has a significant effect. The output uncertainty and rate of the process had less significant effects on the scores,

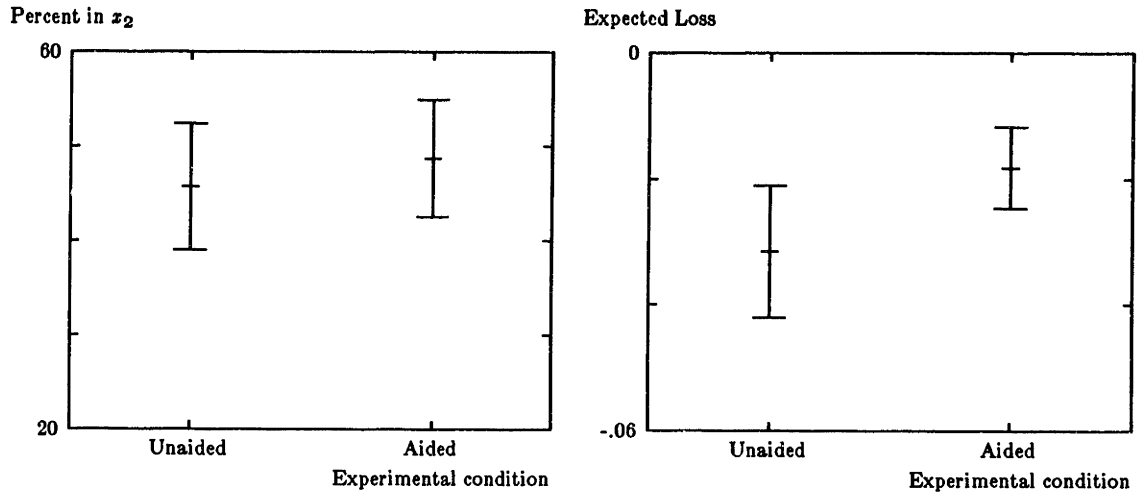


Figure 4.2. The effect of the decision aid on decision making performance using two different scoring rules.

Parameter	Value	Subject 1		Subject 2		Subject 3		F Samp	F Crit
		μ	σ	μ	σ	μ	σ		
Aid	No	48.5	13.6	45.5	10.4	40.6	9.3	7.78*	6.85
	Yes	51.7	10.8	47.8	12.4	46.6	12.8		
Process	0.187	53.7	20.1	52.2	14.0	42.0	12.3	6.27*	4.79
Uncertainty	0.333	45.7	9.1	43.3	3.3	39.5	8.6		
I_p	0.418	46.2	5.2	40.3	5.7	40.3	5.7		
Output	0.131	53.9	7.8					3.79	3.95
Uncertainty	0.391	47.8	14.7	45.4	9.9	40.4	9.5		
I_q	0.651	48.0	10.9	47.2	12.7	41.7	10.0		
	0.911	44.5	17.4	43.6	7.5	39.7	8.3		
Process	0.133	43.9	4.7	41.6	6.7	44.1	6.1	2.61	3.48
Rate	0.267	43.6	13.4	38.8	5.3	40.1	6.8		
	0.533	55.0	17.9	48.6	10.6	41.4	11.7		
	1.067	51.7	11.0	50.9	11.8	36.8	10.1		
	2.117			47.0	9.8				

Figure 4.3. Summary of results and analysis of variance. A star (*) indicates statistical significance at the 0.05 level.

but clear trends can be seen in the data for these parameters as well.

Figure 4.4 shows the effect of the decision aid as a function of the process

parameters I_p and I_q . We see that the effect of the decision aid as a function of the process uncertainty is weak. The effect as a function of output uncertainty is strong, however, and the decision aid seems to be most helpful when the output uncertainty is high. The net effect of the decision aid seems to be to bring the decision making to a moderately good level that is somewhat independent of output uncertainty.

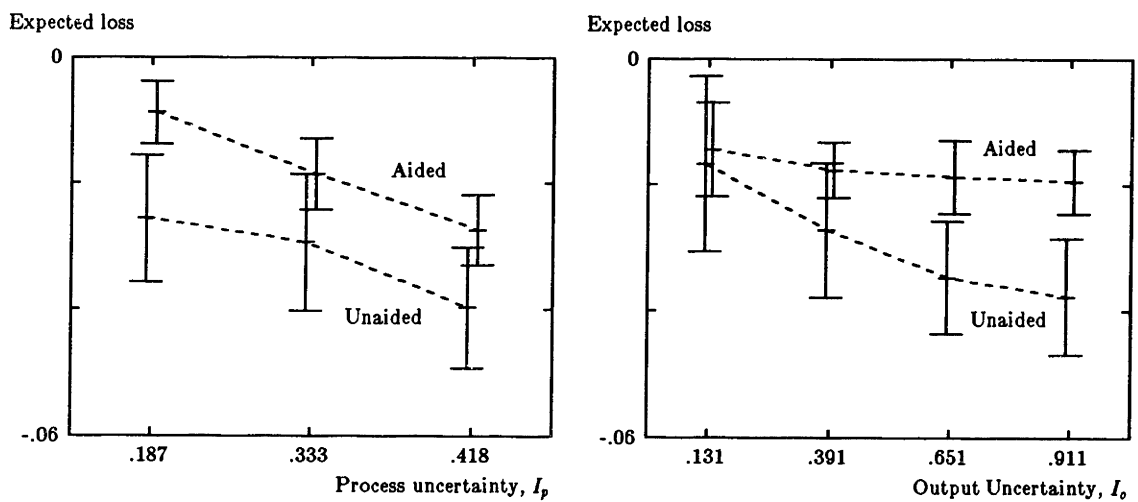


Figure 4.4. Effect of decision aid as a function of process parameters.

The effect of process rate on the effectiveness of the decision aid is shown in Figure 4.5. There we see that the decision aid is most helpful for higher decision rates. This shows that the subject is more pressed for time executing his own decision algorithm than when he is using the decision aid.

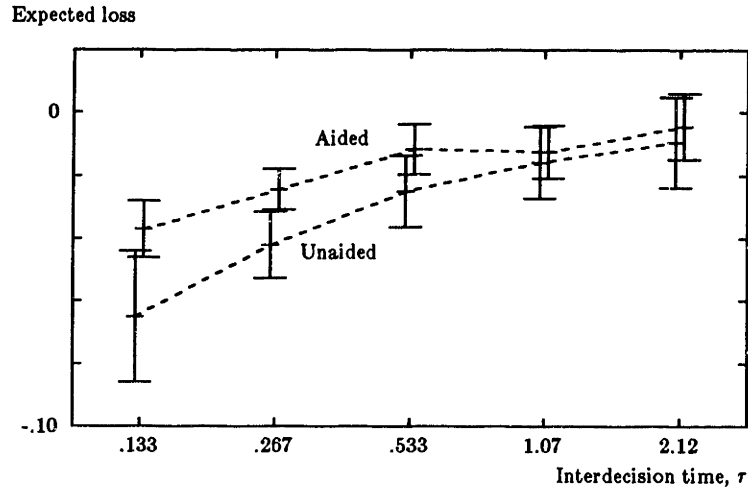


Figure 4.5. Effect of decision aid as a function of process rate.

We define gross errors for this task as any decision made with an expected loss of greater than 0.240 and slight errors as decisions having an expected loss of greater than 0.030. This resulted in nominal error rates of five percent and twenty percent respectively. The error rates as a function of process and output parameters is given in Figure 4.6.

The use of slight errors as a performance measure produced nominally equivalent results as the use of expected loss. For all gross errors, and for all output uncertainty conditions, the decision aid produced flat curves. This suggests presence of the aid in these conditions reduces error levels to some nominal level. The aid seems to correct for the human's biases resulting from high output uncertainties.

4.2.6 Models Predicting Detailed Actions

The observables presented to the subject in an experiment can be input to a model of the subject's decision making to predict his action during each decision

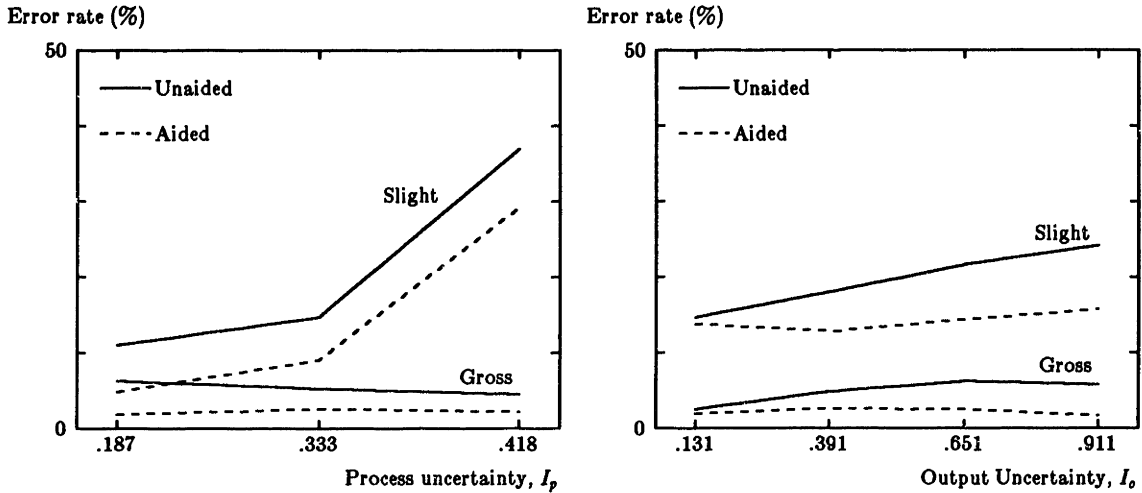


Figure 4.6. Gross errors and slight errors as a function of process parameters.

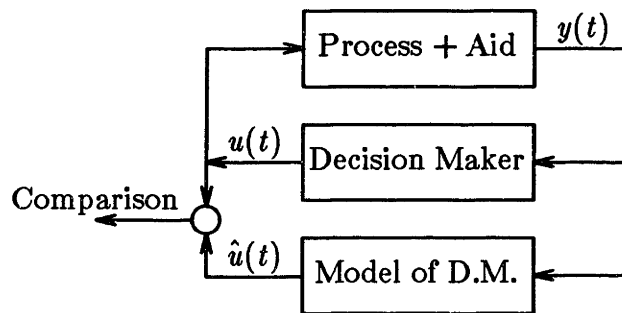


Figure 4.7. A model of the human may be used to predict his detailed action behavior.

interval. This is depicted in Figure 4.7 . Clearly a perfect model will predict the human's actions correctly 100 percent of the time.

Five models of the human subjects were examined for this experiment. Four were formed by combining one of two state estimation models with one of two action selection models. The state estimation models were termed the "Bayesian" or normative model, and "iconic memory" model in which an average of the recent observations is taken as the belief state. The action selection models were

called "optimal" which is normative, and "truncated state" in which the maximum value in the state estimate is taken as the state before applying the selection rule. The fifth model, called the "simple" model, was a thoughtless rote procedure where each observable maps directly to an input without computing a belief state.

Model Number	Model Name	State Estimate	Control Rule	Subject Number		
				1	2	3
1	Simple	-	-	81.0	72.7	62.4
2	Optimal Bayesian	Bayesian	Optimal	59.8	51.9	56.3
3	Truncated Bayesian	Bayesian	Truncated State	82.5	76.3	65.3
4	Optimal Iconic	Iconic Memory	Optimal	80.4	70.4	61.3
5	Truncated Iconic	Iconic Memory	Truncated State	81.0	73.5	63.1

Figure 4.8. Ability of various models to predict actions by each subject.

The results of these comparisons are given in Figure 4.8 . There one can see that the worst predictor of detailed actions is the completely normative model, while the simple model does a good job of predicting actions. None of the models predicts correctly greater than 82.5 percent of the time.

4.3 Experiment Two: Trapping Task with Negative Evidence of State

4.3.1 Qualitative Process Characteristics

Like the previous experiment, this experiment used a process having three states, three outputs, and two actions. Under action u_1 the process state changed frequently, while under action u_2 the state changed only occasionally. The

observable in this experiment was wrong more than it was right. This was to force the subjects to hold an explicit belief state. On some trials, a decision aid was available as before.

The cover story was the same as before, except for the description of the process characteristics. "When ice is put under each bucket, the frog jumps little. When fire is put under each bucket, the frog jumps frequently. Most of the time the dog is wrong."

4.3.2 Quantitative Process Characteristics

The goal in this experiment was to consider the use of the decision aid in greater detail. Therefore decision rate was held constant at 1.0 seconds, and the process characteristics were held constant throughout. They were,

$$\Phi(u_1) = \begin{bmatrix} .96 & .04 & .02 \\ .02 & .92 & .02 \\ .02 & .04 & .96 \end{bmatrix} \quad \Phi(u_2) = \begin{bmatrix} .5 & .5 & .25 \\ .25 & 0 & .25 \\ .25 & .5 & .5 \end{bmatrix}$$

$$\Psi = \begin{bmatrix} .1 & .45 & .45 \\ .45 & .1 & .45 \\ .45 & .45 & .1 \end{bmatrix}$$

$$\mathbf{r}^T = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

4.3.3 Experimental Design

Six subjects who were undergraduates in engineering participated in the experiment. They were given a qualitative description of the experiment as in Section 4.3.1, and the quantitative description as in Section 4.3.2 .

There were four conditions used in the experiment. They were 1) visible state, 2) outputs only, 3) outputs and aid, and 4) aid only. The visible state condition was used as a control and was useful for establishing a base rate for unintentional errors. The outputs alone condition represents the system before a decision aid is added. The third condition represents the normal case when the decision aid is available. The aid only condition ensures that the subject will be trying to use the decision aid for his decisions since it is the only source of information about the process state.

Four trials of sixty decisions each were run per subject in each of the conditions. The trials were grouped in pairs in the order visible state, raw only, raw plus aid, aid only, and then in pairs in reverse order. All four trials of a single condition were tested successively, with a short rest period between each. Subjects were given training in each condition consisting of two trials of sixty decisions plus additional trials until they felt comfortable before starting any recorded trials. Typically, the training session would last forty five minutes, with all training occurring before any measured trials were performed.

4.3.4 Usefulness of Decision Aid

Figure 4.9 shows a summary of the subjects' scores by experimental condition. The basic results are as we would expect. First, the rate of errors in the control condition was low, indicating that the decision time was long enough and the subjects understood the goals of the task. Second, the worst condition was the unaided case followed by the output plus aid condition, while the best was the case where only the decision aid information was available.

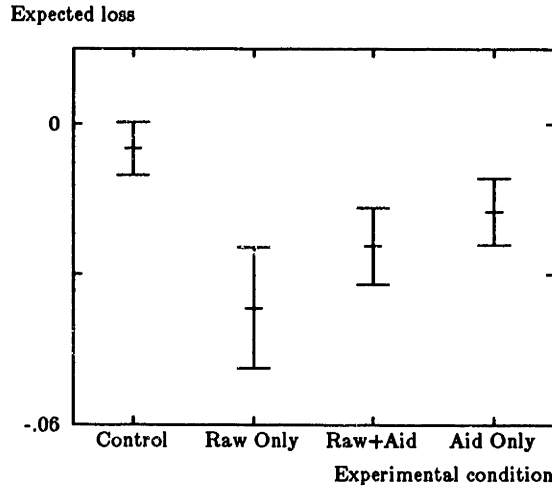


Figure 4.9. Summary of results for all subjects by experimental condition.

4.3.5 Ignoring the Decision Aid and Workload

A more interesting result is obtained by looking at the results on a subject by subject basis. Figure 4.10 shows the results for two of the subjects.




Figure 4.10. Summary of results for two different subjects.

In this summary, one can see that while for one subject his decision performance increased dramatically when both aid and outputs were available, for the other subject his decision performance did not show significant improvement over the

raw observables condition. We conclude that for subject *S2* the aid was used when it was available, while for *S5* the aid was ignored when possible.

This conclusion is supported to a degree by their protocols. Consider a portion of *S2*'s explanation of his control strategy when both the raw state information and the decision aid were available:

I guess having [the decision aid] there gives you a feeling of if you have it or not. If you didn't have the nerd you wouldn't be able to watch the percentages build up... With the dog what I would try to do is just look for a pattern ... And then with the nerd I tried to watch the percentages...

Subject *S1* said he was using the decision aid: when it was available:

I would mostly go with the nerd. But sort of keep an eye on the dog. Let's say there were two that had a pretty high percentage and the dog was on one of them. I would go with the other one.

Not only has this subject described his preference for the decision aid information, he has also discussed how he updates his belief state when both the raw information and the decision aid are available.

Now consider the statements made by *S5*:

With both [the raw outputs and the decision aid], I tended to pay more attention to the dog, because ..., I guess they are pretty similar but, all [the calculations are] going to tell me ... are the likelihood of the frog jumping into the second or third bucket.

Clearly the subject prefers the raw outputs to the decision aid even though his decision making improves when he must use the decision aid.

4.3.6 Extracting a Selection Rule

In Section 2.3.5 an approximate method for determining the optimal selection rule is presented that results in a division of the space of possible belief states

linearly into sections where various actions should be selected. For a system having three possible states, as ours does, the belief state consists of three values on the interval $[0,1]$, two of which are independent. A selection rule can be plotted for visualization on a plane, and the regions corresponding to various actions can be shown.

If the decision aid is the sole source of information as in the fourth condition, one can assume that it corresponds to the belief state of the subject, and the selection rule used can be plotted in a scatter plot. This has been done for subject two in Figure 3.11. The normative selection rule would be a line at $f_1=.2$.




Figure 3.11. Approximation of selection rules for two subjects. An x shows action u_1 was chosen given the belief state while an o shows action u_2 was chosen. The dashed line separates the two regions.

4.4 Experiment Three: Approximation to a Continuous Task

To provide a dynamic process that approximated a process having continuous states, a dynamic process with eighty one states was constructed.

4.4.1 Qualitative Process Characteristics

The process had eighty one possible states, eighty one possible observables, and three possible actions. The state transitions were chosen so that the particle had one of nine possible positions and nine possible velocities. It had the behavior of a particle bouncing between two walls. This is shown in Figure 4.12.




Figure 4.12. The dynamic process in experiment three was given behavior of a particle bouncing between two walls.

Action u_1 was labeled "push left" and produced a force on the particle such that the velocity tended to decrease. Action u_2 was called "leave alone" and neither forced the particle to the left or right. Under action u_3 , labeled "push right," the velocity tended to increase. The state was not directly observable. In this experiment, only normatively derived state information processed in various ways was presented to the subject.

4.4.2 Quantitative Process Characteristics

The time between decisions was fixed at 2 seconds. The state transition matrices $\Phi(u_i)$ and the output matrix Ψ were defined implicitly as follows.

$$\begin{aligned}\phi_{ij}(u_k) = & D(i, R(X(j)+V(j), V(j)+F(k)-1)/4 \\ & + D(i, R(X(j)+V(j), V(j)+F(k))/2 \\ & + D(i, R(X(j)+V(j), V(j)+F(k)+1)/4\end{aligned}$$

where

$$X(j) = j \text{ modulo } 9-4$$

$$V(j) = j/9-4$$

$$F(k) = k-1$$

$$D(i, j) = \begin{cases} 1 & : \quad i=j \\ 0 & : \quad \textit{otherwise} \end{cases}$$

$$R(x, v) = \begin{cases} N(4-x, -L(v)) & : \quad x > 4 \\ N(-4-x, -L(v)) & : \quad x < -4 \\ N(x, L(v)) & : \quad \textit{otherwise} \end{cases}$$

$$L(v) = \begin{cases} 4 & : \quad v > 4 \\ -4 & : \quad v < 4 \\ v & : \quad \textit{otherwise} \end{cases}$$

$$N(x, v) = x+4+9(v+4)$$

To give an interpretation to all these definitions, $X(j)$ is the position associated with j , $V(j)$ is the velocity associated with j , $F(k)$ is the nominal force or change

in velocity from u_k , $L(v)$ is v limited to the range $[-4,4]$, $N(x,v)$ is the index number associated with x and v , and $R(x,v)$ is the index associated with x and v after limiting and reflections.

The reward function decreased linearly with distance from the center position, with no dependence on the "velocity." It is defined as,

$$r_i = 1 - |X(i)| / 4$$

4.4.3 Experimental Design

The goal of the experiment was to determine how much information is used by the subjects when they are presented a display consisting of normatively derived state information including probability information. Six different conditions were tested in the experiment. In each condition different information was displayed about the process state. In each case the information was normatively derived state information simplified in various ways. The six displays used are shown in Figure 4.13.

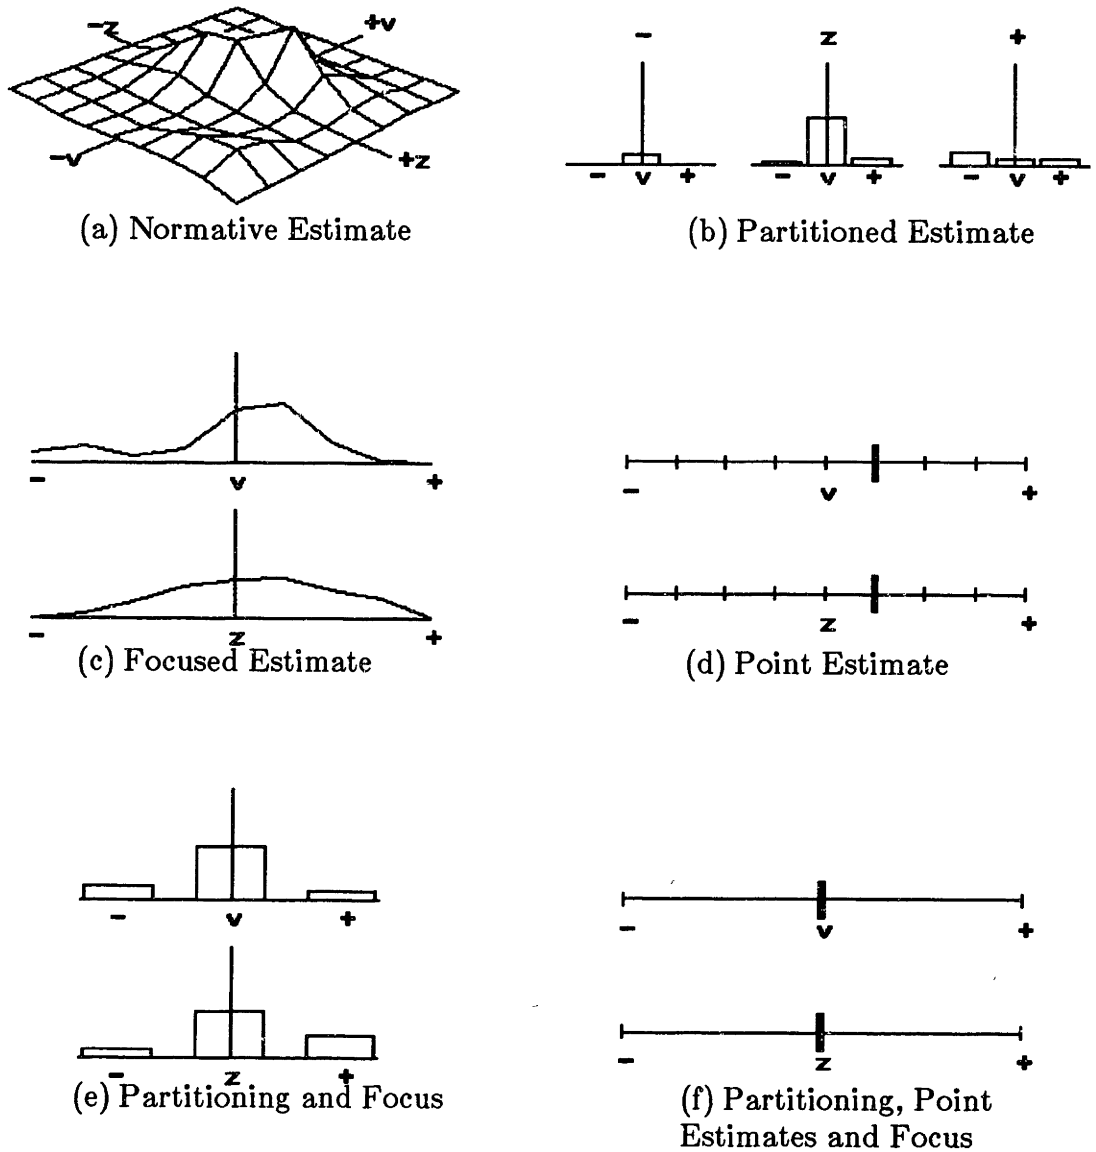


Figure 4.13. Six different displays of normatively derived state information were used in experiment three.

Two undergraduates in engineering participated in the experiment. The experiment consisted of four trials in each condition, and in each trial sixty successive decisions were made. The state, output, and decision were recorded each time to allow reconstruction of the displays in later analysis. The decision time was constant at 2 seconds throughout the experiment, and the time remaining for

a decision was displayed on the screen. This is shown in Figure 4.14.




Figure 4.14. Sample of display showing decision time remaining.

The subject could select his action at any time during the decision period by pushing an arrow key on the keyboard. At the end of each decision period, the last arrow key pressed by the subject was recorded as the decision. If no arrow key was pressed the "leave alone" selection was entered as the decision.

4.4.4 Approximations to Subjects' Selection Rule

As in experiment two, an approximation to the subject's selection rule was determined from the data. Since the belief state consists of a vector of eighty one values, eighty of them being independent, a simplified representation of the belief state must be used for our visualization. The process dynamics were specifically chosen to represent a particle having mass, so we use a phase space representation

of the mode of the belief state to plot points in the belief state in two dimensions.

Since the subject was only given normatively derived state information, we will assume that his belief state is identical to what is presented. When this assumption is made, we can plot the subject's presumed belief state as the dependent variable and use the subject's action as the independent variable. This has been done for one subject in Figure 4.15. This provides a quick way to visualize how consistent the subject is, how many mistakes he made, and what his approximate selection rule looks like.

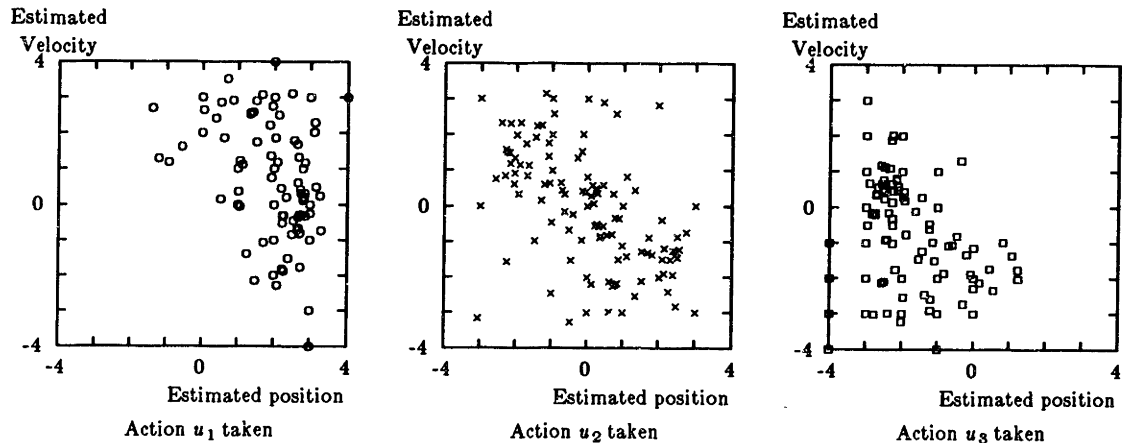


Figure 4.15. A scatter plot of the presumed belief states leading to each possible action.

4.4.5 Models Predicting Detailed Actions

As before, models of the human's decision making were used to predict the detailed actions taken by the subject. These models were made up of various models of the data simplification method, utility function, and selection rule. The six models for data simplification that were used correspond to the six displays used in the experiment. Three models for the utility function were used, and this

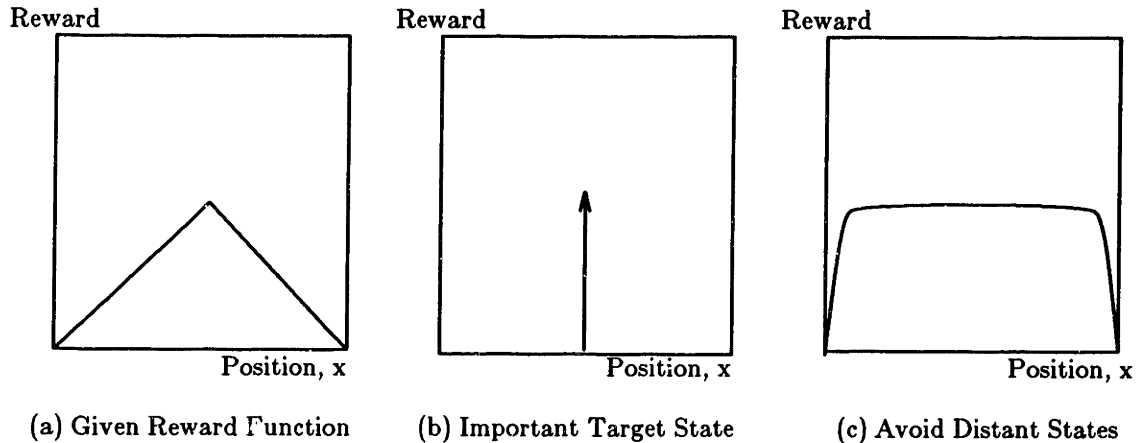


Figure 4.16. Three possible reward functions that the subjects could be using.

is in Figure 4.16. The first is the utility function as it was given. The second is a simplified utility function in which the operator concentrates only on getting to a target state. The third is a function such that the operator will avoid certain states. Two selection rules were used. One is the short sighted decision rule of Section 2.3.4 and the other is the normative rule of section 2.3.5.

The results of the predictions are given in Figure 4.17

Model Number	Decision Rule #1			Decision Rule #2		
	Util Func	Util Func	Util Func	Util Func	Util Func	Util Func
	#1	#2	#3	#1	#2	#3
0	42.5	43.6	41.7	40.6	41.4	41.1
1	51.1	51.9	50.3	49.4	57.2	55.6
2	42.2	43.6	39.7	40.6	40.8	40.8
3	50.8	51.9	50.0	48.6	57.2	55.6
4	38.0	37.8	39.2	37.2	21.9	41.1
5	51.7	52.5	50.8	49.1	57.2	55.6

Figure 4.17. Predictability of various models of decision maker.

In it we see that the utility function model has a slight effect on predictability, with the target state utility function providing better predictions of human

decisions. The selection rule used had a varying effect on model predictability. The simplification rule had a much higher effect, with models 1, 3, and 5 scoring virtually the same.

An *ad hoc* model of one subject was made by combining simplification model 2 with a selection rule derived from his belief state scatter plots. This model yielded much higher predictability than any of the previous models as shown in Figure 4.18.

Experimental Condition	Predictability (percent)
0	63.3
1	75.0
2	76.7
3	81.7
4	73.3
5	73.3
Overall:	69.4

Figure 4.18. Predictability of *ad hoc* model of subject one.

4.5 Simplification Models of the Human

From the previous experiments, some descriptive models of human decision making for dynamic decision making tasks have been developed. Because they all involve methods or computations that appear to simplify the computational burden on the decision maker, we call them simplification models.

4.5.1 General Basis for Forming Models

The basis for the descriptive models of the next three sections is the normative model of decision making presented in section 2.3. The human decision

maker is viewed as carrying out an imperfect version of this normative algorithm. Several components were identified in the normative model including a belief state, a utility function, a prediction process, an updating process, and a selection rule. Additionally, the prediction process and the selection rule require a model of the process dynamics, while the updating process requires a model of the process outputs. For the decision maker to be normative, all these elements have to be normative. The models must be accurate, the belief state complete, and the processes perfect.

Thus many models of the human can be formed by beginning with purely normative components and making one or more components subnormative in various ways. The following sections present some descriptive decision making models that partially explain the observed decision making behavior.

4.5.2 Substituting Rote Procedures for Contemplated Action

One way in which a human decision maker may simplify his cognitive task is to substitute a simple rote procedure for the laborious normative procedure. An example of this type of simplification is shown in cartoon form in Figure 4.19.




Figure 4.19. By substituting a rote procedure for a normative one, an operator saves mental effort at the expense of performance.

In certain decision situations, a rote procedure may be substituted with negligible effect on overall decision making performance, though in general there will always be some effect. In experiment one, a simple rote procedure model of the human was as good as any other in predicting the specific actions taken, yet their performance at the task was only marginally reduced by the substitution. We say rote procedure because it does not involve the computation of a belief state with the passage of time and presentation of new state evidence, but rather blindly operates on incoming data in a predetermined stimulus→response way.

Good use of rote procedure substitution requires premeditation about the process behavior. In many real systems operators are asked to follow procedural handbooks for certain conditions and circumstances rather than think deeply about the process state and carefully select an action. Presumably the considered

use of procedure handbooks results in better average decision making by the operator. Sometimes an operator may need to respond to process states that are unfamiliar to him or rarely visited, so he will not have a practiced way of thinking about the situation. In these cases a procedural handbook may speed him to an action that is good, and may be preferable to having him cogitate ponderously while the unattended problem runs out of control. Precisely when it is desirable to have the human applying his unique reasoning abilities is still an open question.

4.5.3 Partitioning, Making Point Estimates, and Focusing

If an operator is presented a normatively derived state estimate that consists of a probability distribution over all possible states, he will almost certainly find it is too much information to comprehend. In experiment two of Section 4.3.5 some subjects were already uncomfortable using all the information presented, even for a system having only three possible states. Many realistic systems have several continuous variables mixed with discrete variables, and the presentation of information for these systems presents extraordinary difficulties.

Even if the information can be presented, it is another problem for the human to comprehend what is presented. There is evidence that he only uses the barest of details from a distribution even when presented with the entire distribution, and he may unwittingly simplify the information in a dangerous way. The terms partitioning, point estimates, and focus will refer to three ways in which the human can simplify his presented distributions to cope with the excessive information. While Figure 4.13 shows a distribution altered by these operations in

various formats, Figure 4.20 shows a complete distribution before and after the operations of partitioning, point estimation, and focus in a single format.




Figure 4.20. A distribution before and after simplifications.

Partitioning. A variable that can take on many different values in the true process is considered as having only a few possible values for the decision maker. An example may be a certain temperature measurement in a system. Instead of making a reading of 79.5 F and then considering all its ramifications, the operator may directly translate the reading into "Too high," "About right," or "Too low." Within one of these superstates, the operator does not distinguish between different temperatures. Therefore his inferences and conclusions may be exactly the same for slightly different true temperatures.

Making point estimates. A single point is selected from a presented distribution and all reasoning proceeds as if the process state were known to be that state exactly. This reduces the amount of information from a distribution over a real valued variable to a single real valued variable. As with any simplification, this can lead to systematic forms of dangerous behavior. We will take this issue up in Section 4.6. Failure analysis presumes the validity of this type of reasoning.

Focusing. This is considering the information in a single variable when a distribution is presented and ignoring what information is contained in the joint probability distribution. When the variables have a high degree of dependency, information will be removed by this simplification.

4.5.4 Using Simplified Task Characteristics

Another way in which an operator can simplify his job is to misunderstand the process characteristics as being those of a system with which he is already familiar. This will be especially true when an operator is working with a system that is new and unfamiliar.

Section 3.4 presented an experiment that was like a tracking task with which many people are familiar. The measured selection rules for one subject were presented in Figure 4.15. Here we saw well-differentiated regions of actions. But consider what the scatter plot for a normative decision maker would look like as presented in Figure 4.21.

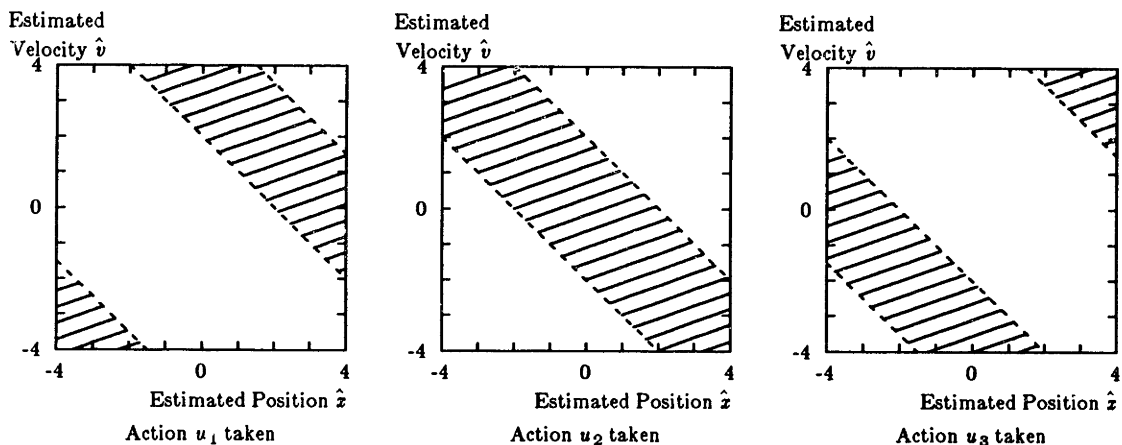


Figure 4.21. Scatter plot of ideal decision maker in experiment 2.

There are regions in the corners of the bounded phase space of the ideal decision maker that do not show up in the subjects. This is because the specification of the walls of the process as bouncy. When moving toward the wall, the best action is to force the particle toward the wall so it will rebound near the desired state, not away from it as the subject consistently did.

Now consider the scatter plot of Figure 4.22.




Figure 4.22. Scatter plot of ideal decision maker for particle without walls.

This is the scatter plot for an ideal decision maker controlling a particle in open space. This is almost identical to the experimentally produced scatter plot of experiment three. Thus we conclude that the subject used a convenient and familiar process model, a particle without walls, instead of the more accurate particle between bouncy walls.

4.6 Effects of Simplifications on an Ideal Decision Maker

Given that human operators simplify incoming data and internal represen-

tations, what is the theoretical effect of these simplifications on decision making? In the remainder of this chapter we will try to examine this question as it relates to presented experiments and models.

4.6.1 Effects of Rote Procedure Substitution

As we said, a rote procedure is a predetermined stimulus response pattern without any belief state. Since there is no belief state, the only states of the process plus operator system are those of the process. When the controlled process has a finite number of possible states, the system behavior is itself a finite state process and the expected system output can be solved explicitly.

In experiment one, a rote procedure that was a good predictor of decision making was,

$$u(t) = \begin{cases} u_1 : & y(t)=y_3 \\ u_2 : & y(t)=y_1 \\ u(t-dt) : & y(t)=y_2 \end{cases}$$

This rote procedure does have an independent state, that of the last action. The system state vector will be augmented to include the outputs and this decision maker's state to be

$$x_{aug} = \begin{bmatrix} x \mathcal{E} y_1 \mathcal{E} u_1 \\ x \mathcal{E} y_2 \mathcal{E} u_1 \\ x \mathcal{E} y_2 \mathcal{E} u_2 \\ x \mathcal{E} y_3 \mathcal{E} u_2 \end{bmatrix}$$

The system transition function may be written as,

$$\Phi_{aug} = \begin{bmatrix} \psi_1 \Phi_1 & \psi_2 \Phi_1 & \psi_2 \Phi_1 & \psi_3 \Phi_1 \\ \psi_1 \Phi_1 & \psi_2 \Phi_1 & \psi_2 \Phi_1 & \psi_3 \Phi_1 \\ \psi_1 \Phi_2 & \psi_2 \Phi_2 & \psi_2 \Phi_2 & \psi_3 \Phi_2 \\ \psi_1 \Phi_2 & \psi_2 \Phi_2 & \psi_2 \Phi_2 & \psi_3 \Phi_2 \end{bmatrix}$$

The reward function is given by,

$$\mathbf{r}_{aug}^T = \begin{bmatrix} \mathbf{r}^T & \mathbf{r}^T & \mathbf{r}^T & \mathbf{r}^T \end{bmatrix}$$

4.6.2 Effects of Point Estimates and Partitioning

In the cases of point estimate, partitioning, and focus simplifications, it is more difficult to analyze the effects of these on decision making performance analytically. There is a fair amount of literature devoted to "fault diagnosis," [1][2][60][61][34] an activity which is directed toward determining the single cause (point estimate) of unexpected process behavior, so it is important to know what effects such a view has on decision making. We consider two hypothetical examples in which these simplifications lead to poor decision making performance.

In Figure 4.23, an example is illustrated in which making the point estimate simplification leads to an incorrect decision. The real belief states and utility function are given on the top. There, the probable future state of the process under actions u_a and u_b are shown. The (adjusted) utility function defined over the states is also shown. By taking the expected value of each decision, one can see that action u_b has a higher expected value. Superimposed, the belief states and utility function are shown after the point estimate simplification has been made. Here the subject will pick the incorrect action u_a .

Figure 4.24 gives an example in which partitioning leads to incorrect decisions. Again the real unpartitioned process appears on the left and u_b should be




Figure 4.23. Hypothetical example showing point estimate simplification that leads to subnormative decisions.

chosen. But when the system is partitioned into one having fewer possible states as on the right, u_a appears to be the best action.

4.6.3 Effects of Task Simplification

In Section 4.5.4, it was discussed how a simplified or familiar set of task characteristics may be substituted for the true characteristics. Since there is little understood about the way humans represent system behavior in general, it is difficult to apply analytical techniques for these simplifications as well.

Original State Variables

Partitioned Variables




Figure 4.24. Hypothetical example showing partitioning simplification that leads to subnormative decisions.



The Libraries
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Institute Archives and Special Collections
Room 14N-118
(617) 253-5688

This is the most complete text of the
thesis available. The following page(s)
were not included in the copy of the
thesis deposited in the Institute Archives
by the author:

pg 102

5. Human Prediction of Naturally Occurring Rotational Motion

An experiment was performed to determine humans' natural abilities to observe and predict three dimensional rotational motion. While simple spinning motion is a familiar sight on earth, tumbling motion of a rotating body is commonly observed in space and seldom observed in everyday life.

Naturally occurring rotational motion of rigid bodies cannot be described by linear differential equations. Depending on the body's mass parameters and initial conditions, motion can be classified as simple spinning, nutation, or tumbling. The types of motion require, respectively, one, two, and three varying state parameters to make predictions about future states. We also refer to these separate cases as having one, two, or three degrees of freedom.

5.1 Experimental Design

5.1.1 Description of Experimental Task

The experimental description is as follows: The subject is seated at a computer generated display. After a brief (2-4 sec) period in which he sees only a blank screen, a rotating rectangular object appears on the screen for a randomly selected time ($t_o=0.5-32.0$ sec) which is called the observation interval. The object is marked so it can be unambiguously oriented in space.

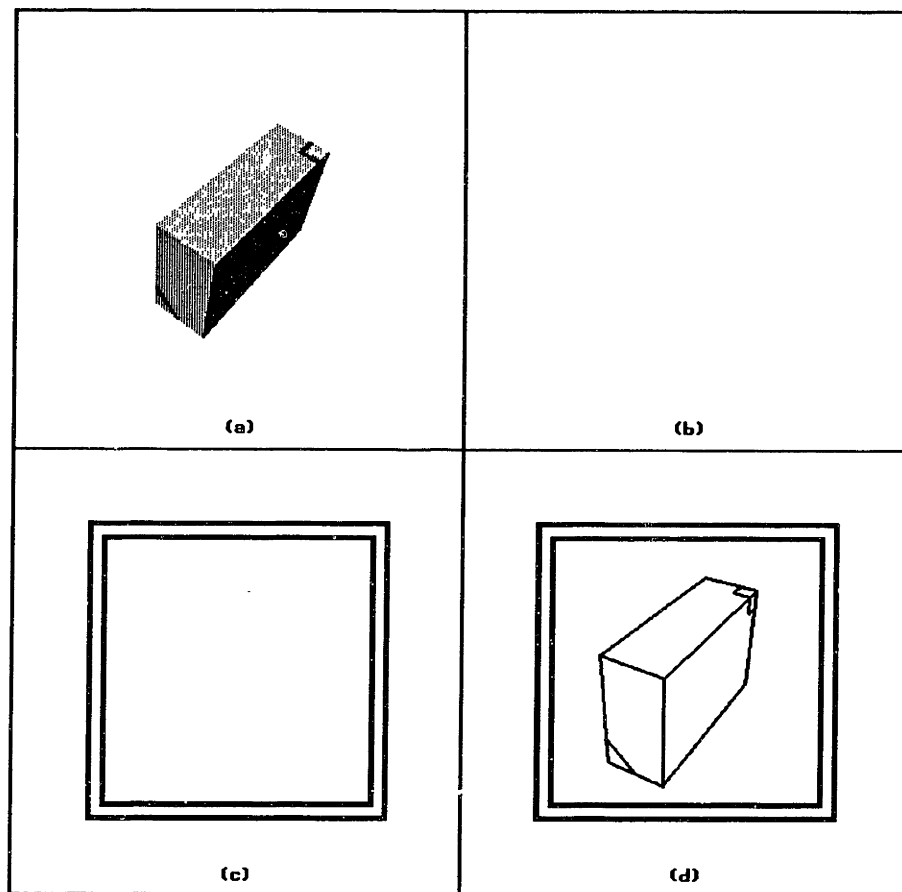


Figure 5.1. The four parts of a single experimental trial.

Figure 5.1(a) shows a frame from this portion of the trial, while 5.1(b-d) show successive portions. In all trials the object was rectangular with sides having

ratios 1-2-3. The mass parameters used in simulating the motion were chosen to reflect this geometry, though in some trials they were mis-aligned with the visual geometry. The screen then goes blank for a randomly determined time ($t_p=0-32$ sec) called the prediction interval, as in 5.1(b).

The subject is instructed to imagine the position of the object and track its original motion during this time. At the end of this period a prompt shown in 5.1(c) appears. The goal of the experiment is to indicate as closely as possible where the object would have been at the moment the prompt first appeared. In Figure 5.1(d), the subject activates a three dimensional cursor to indicate his estimate of the object's position. The error in his estimate is recorded by the computer and the next trial is presented. The calculated error is the angle through which the estimated position must be rotated to bring it into alignment with the true position at the time the prompt first appears.

5.1.2 Description of Important Task Characteristics

We wish to characterize human ability to predict naturally occurring rotational motion as a function of certain parameters of the motion. The parameters that we have used to decompose the motion are:

Degrees of Freedom The initial conditions are set so that the resulting motion requires one, two, or three varying parameters (state variables) to be fully described.

Observation Interval The period during which the human may observe the system undergoing its natural motion. One would

expect longer observation intervals to produce better performance.

Prediction Interval The period of time over which the human must predict the object's motion. In the absence of additional observations, one would expect the errors to increase with increasing prediction intervals, and after some period, the prediction will be no better than chance.


Rotation Rate The rate at which the object is spinning. One would expect fast rates to increase errors, but experience in other domains has shown that extremely slow rates may also cause difficulties.

Mass Alignment Whether the mass distribution of the object matches the visual geometry. This was varied to yield information about people's natural intuition regarding rotational motion.


Screen Alignment Whether the rotational axis of the body is aligned with the plane of the display. This is well defined only for 1-DOF motion. Other studies have shown that rotation in depth is processed identically to rotation in the plane of the screen for 3-D objects.

Figure 5.2 shows several sequences intended to give the reader a feel for how


the degrees of freedom parameter relates to observed dynamic motion of the body. In each frame several positions have been superimposed to give a better understanding of the resulting motion as it might be perceived in the experiment. Many other combinations were used in the experiment that are not shown in the illustration.



(a) Simple spin



(b) Nutation



(c) Tumbling

Figure 5.2. Several examples of motion and its decomposition into degrees of freedom.

5.1.3 Trials, Order, Etc.

For each subject, the experiment consisted of either 132 trials, or fifty minutes of trials, whichever came first. The 132 trials were randomly ordered and broken down as follows: Three "primary groups" of 41 trials were selected,

one for each degree of freedom. A primary group of trials consisted of one each of trials having prediction interval and rotation rate coming from Figure 5.3. To these 123 trials were added six 1-D.O.F. screen-aligned trials, with other parameters randomly selected.

		Prediction interval							
		0	.5	1	2	4	8	16	32
Rotation rate (rad/ sec)	.034	x	x	x	x	x	x	x	x
	.068	x	x	x	x	x	x	x	x
	.125	x	x	x	x	x	x	x	
	.25	x	x	x	x	x	x		
	.5	x	x	x	x	x			
	1	x	x	x	x				
2	x	x	x						

Figure 5.3. Rotation rates prediction intervals for a primary group.

5.1.4 Experimental Subjects

Fifteen engineering school students and graduates were used as subjects in the experiment. Nine of these subjects had greater than three years of experience with either gymnastics or diving, and an average of four years experience. These subjects were chosen with the assumption that their experience in these areas has given them additional familiarity with rotational motion. The other seven subjects were not experienced in either gymnastics or diving.

5.2 Data Reduction Method

5.2.1 Model of Human Performing Experiment

The model of human performance in this task is that the error of the human's estimate will increase with time in the absence of observations and decrease in time with additional observation. Refinements to this model may be

made by assuming a specific model for the growth of errors; an exponential growth model can be used as a starting point, and it has a strong basis in linear systems theory. The errors are assumed to come from two sources - a finite precision position representation and imprecision in motion transformation operations.

Related work has shown that people do mental transformations on visually stored images of external objects to track rotation [42]. If this is true, then there must be some mechanism to deal with various rotation rates. Either a fixed frame-rate can be used by the human image processing system or the frame-rate can be varied with a fixed transformation being used. Each system would of course result in different error characteristics. Even if the human uses continuous representations or transformations, there will be some nominal level of noise generated by their use, and the value of these levels has important practical significance.

5.2.2 Bayesian Estimation of Model Parameters

To analyze the data, a method proposed and developed by Max Mendel was used [38]. A sketch of this method appears in Appendix F. The data were used to estimate the growth rate of squared angular error in orientation. This will correspond to observing the growth rate of the variance of the sample data. By examining the posterior distribution over the model parameter, we see our own uncertainty about our parameter estimate, *i.e.* the variance of the variance. Classical methods are not able to provide this type of information.

5.3 Preliminary Results

5.3.1 Removal of Outliers

Though the displayed body was designed to suggest a simple mass distribution that is unambiguously orientable in space, frequently estimates by subjects were approximately π radians out of alignment with the true final orientation. These are almost certainly caused by confusion resulting from the high degree of symmetry in the body. Because these outliers do not reflect inability of humans at the primary task but rather mental slips, an effort was made to remove them from the data. Data generated in this way are from a different model and so may not be considered exchangeable.




Figure 5.4. Histogram of error scores for entire group.

Figure 5.4 shows a histogram of the number of times an error occurred during the experiment. A line representing a distribution over the next data point is also plotted. We can see that for small errors representing most of the data, the data and the estimated model agree well. For large errors, however, there are far more data points than expected, presumably because of non-exchangeable data.

It is well known that the presence of such outliers can corrupt classical methods of data analysis, since data that includes outliers will appear to have much larger statistical fluctuation than data in which outliers have been successfully removed. An example of this that is presented as further justification for removing these non-exchangeable data from the sample is given in Figure 5.5. In part (a), data is reduced that includes outliers, while in part (b), data beyond a certain reasonable level have been removed. The data in part (b) clearly show expected trends that will be presented later in this paper, while the data in part (a) appear to contain excessive noise. Choosing a cutoff level that is too low as in (c) will of course result in the removal of exchangeable data. In later analysis, a cutoff level of $x_i=7rad^2$ has been chosen.

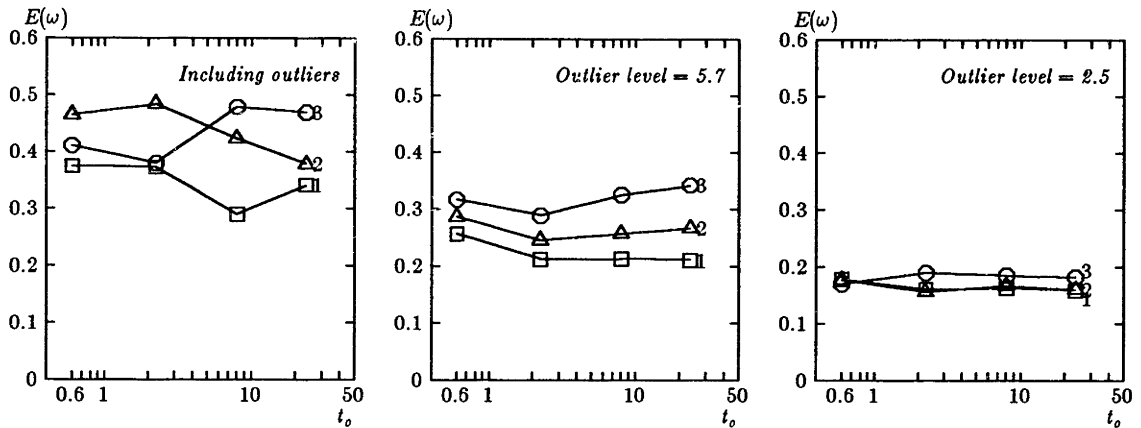


Figure 5.5. The effect of not removing (a) and removing (b,c) outliers on the statistical variation of the reduced data.

5.3.2 Effects of Learning

The effects of learning in the experiment can also be easily examined using the plot of figure 5.6 which shows the posterior spread parameter as a function of trial number after outliers have been removed. From this plot it can be seen that

there is little effect of learning on the subjects' performance. This shows that the short training period provided in the experiment is satisfactory to establish steady state performance.




Figure 5.6. The effect of trial number on performance.

5.3.3 Noise from Input Device

When the subject indicates his estimate of the final position using the three dimensional cursor, there will be some error introduced due only to the man-machine interface. The error due to making this measurement can be determined by looking at the limiting posterior expected error for short prediction intervals and slow rotation rates. Since the experiment was designed to include cases with $t_p=0$ sec, this error can be determined.

Figure 5.7 shows the posterior expected error, $E(x_0)$ as a function of the prediction interval t_p for slow rotation rates, that is $\omega_0 \leq .25 \text{ rad/sec}$. From this plot it can be seen that the data does have redundant error as the prediction interval goes to zero. An approximate value for the limiting error is 0.15 rad^2 . In

further analysis it should be kept in mind that this is an error associated with the subject communicating his estimate to the system via the cursor. Any systematic errors above this level can be said to be the result of effects other than the input device.




Figure 5.7. Posterior expected error as a function of the prediction interval for slow rates.

5.4 Major Results and Supporting Data

5.4.1 Complexity of Rotation

In each trial, the initial conditions were chosen to present the subject with motion having either one, two, or three degrees of freedom corresponding to simple spinning, nutation, or tumbling. Figures 5.2(a)-(c) illustrate examples of each of these types of motion. To evaluate this motion parameter as a predictor of task performance, the posterior error estimates are calculated conditional on the number of degrees of freedom of each trial. These calculations are plotted in figure 5.8(a).

These results are as expected. That is, the degrees of freedom of the motion is a good predictor of task difficulty, with single degree of freedom motion producing the best performance.




Figure 5.8. Posterior error (a) and precision (b) as a function of motion degrees of freedom.

5.4.2 Differences in Group Characteristics

Subjects for the experiment were gathered from two different populations: those having extensive gymnastics or diving experience; and those having little gymnastics or diving experience. Presumably such experience should provide the participant with exposure to more complex rotational motion not otherwise experienced in everyday life. This type of exposure would sharpen the intuition and predictive abilities to the extent that such sharpening is possible. Since we are comparing two groups of subjects, we normally could expect some differences in their performance levels. The presence of one degree of freedom trials in this experiment allow us to remove these differences from the data. Because both groups can be expected to have extensive experience with simple spinning motion,

we can say that both groups have similar backgrounds for this task, and whatever differences appear in the data are intergroup differences.

What we see in Figure 5.8(a) is an intergroup difference that seems to be independent of the stimulus degrees of freedom. The significance of these differences can be compared by plotting the posterior distribution over the spread parameter ω as in 5.8(b). Doing so shows that for the data collected, the differences between groups do not appear to be significant; however those between degrees of freedom do appear significant. Thus the data confirm the model, which says that humans will have more difficulty with more complex motion, but also tends to deny the idea that additional experience with rotating objects will significantly improve performance in the experimental task.

5.4.3 Observation Interval

An important part of the experiment is determining how performance is related to the observation interval, the prediction interval, and the rotation rate of the presented stimulus. These effects are interesting because they may relate directly to parameters of the human information processing system. Figure 5.9(a) shows the expected errors as a function of the observation interval, t_o . The plot is flat over the range of values tested. For convenience, the data has also been plotted in two groups: the gymnasts, and the control group showing the level of fluctuation present after the data has been reduced. In addition, the data has been broken down according to degrees of freedom in figure 5.5(b).

From these plots we must reach the surprising conclusion that the task per-

formance is *independent* of the observation interval, t_o . What this means in practice is that the tested range did not span the intervals that are important in terms of the human information processing system. Apparently, whatever position and velocity estimation will be done by the human for this task are done within 0.5 sec of observation, and that further observation is not used by the subjects in any task-related way.




Figure 5.9. Error levels as a function of the observation interval.

5.4.4 Prediction Interval, and Rate

The design of the experiment requires that the prediction interval, t_p and the rotation rate, ω_0 be considered together. In addition, it may prove enlightening to consider the prediction angle α_p , defined as the product of t_p and ω_0 , and its function in predicting performance levels.

By looking at various combinations of these three variables, it was determined that the prediction angle α_p is a good predictor of performance for a wide range of trial parameters, as shown in Figure 5.10(a). There it can be seen that

with a few exceptions all data lie along a single line, even for different prediction intervals. The exceptions to this are for trials having the slowest rotation rate. For these rotation rates the error appears to be proportional to the prediction interval t_p . This can be seen from Figure 5.10(b).

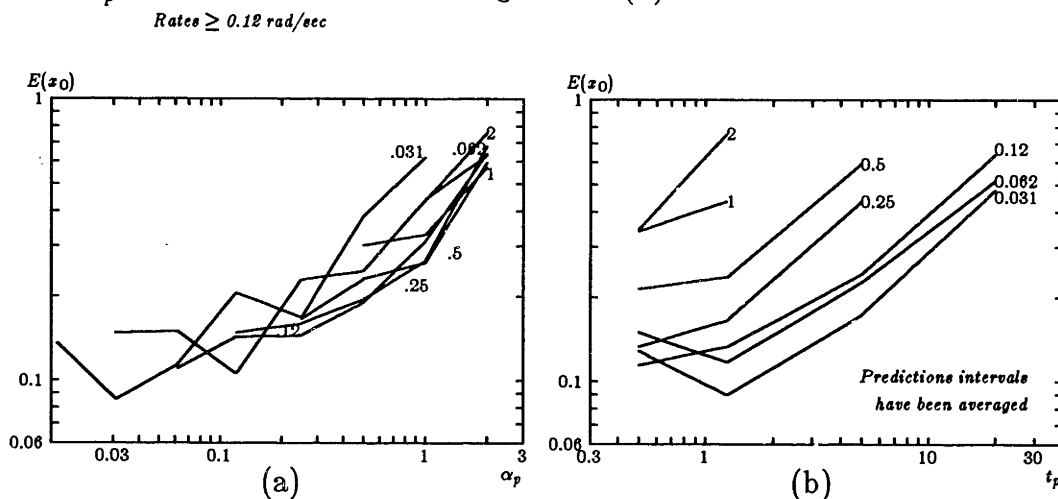


Figure 5.10. Expected error as a function of prediction angle α_p (a) and prediction interval t_p (b).

An explanation of these characteristics could be as follows. The human has at his disposal a homeomorphic representation of the angular position with finite precision and a set of rotation transformations he can apply that also have finite precision. For moderately fast rates, as in most of the experiment, the subject uses a single finite transformation at various rates, so the error will be proportional to the number of times the transformation is applied, or equivalently, proportional to the total angle through which the body rotates. Because of the large amount of scatter in the data, these conclusions must be considered tentative.

5.4.5 Rotation in Screen Depth

The experimental design provided for a small number of trials to be intermixed with normal trials that were one degree of freedom rotation with the

rotation axis perpendicular to the screen. The data for these trials was separated and analyzed, with the results shown in figure 5.11. In figure 5.11(a), the effect of screen alignment for one degree of freedom on the posterior error estimate is shown, with the two and three degree of freedom cases shown only for reference.

When compared to differences based on degrees of freedom, the increase in task performance based on the screen alignment is negligible, and this can also be seen from the posterior precision distributions for one degree of freedom in figure 5.11(b). Note that the data reduction methods allow easy significance comparison for samples of widely varying size and data can be collected until a desired significance level can be achieved. Thus the data support the hypothesis found in the literature that tasks that happen to require only two-dimensional visual processing - as in the screen-aligned cases - are processed no differently from tasks that are visually three dimensional in nature.

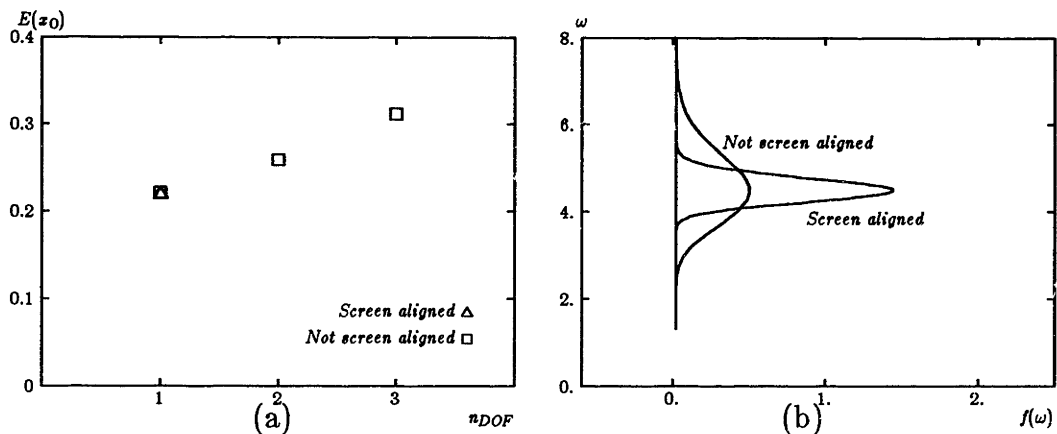


Figure 5.11. Effect of screen alignment on task performance.

5.4.6 Alignment of Inertia Parameters

During testing of the first few subjects, some people reported a sensation

that the stimulus was suddenly accelerating, or looked as though there were moving parts inside. This sensation may legitimately come from two separate causes. First, it is a simple fact that for unforced rotational motion, the angular momentum vector is constant but the angular velocity need not be, and is not for all but simple spinning. Hence for the two and three degree of freedom trials the body will undergo angular accelerations as a normal part of its motion. Second, in everyday life the visual geometry gives important clues to properties of the object for many interactions, so it is reasonable to assume that subjects, if able, will guess about the inertial properties of an object from its visual geometry. Since in some trials the visual geometry gave false indications about the mass distribution of the body, it is conceivable that a sensitive subject would judge the observed motion as surprising, unnatural, or looking like there were moving parts inside.

As a result of these considerations, later subjects were asked to indicate whether they had this subjective experience at the end of each trial. These data were then compared with the true degrees of freedom and alignment to determine the basis of these judgements. The results of this comparison appear in figure 5.12(a). From this data we see that subjects were more likely to indicate surprising behavior as the alignment was degraded, and they were also more likely to do so for two and three degree of freedom trials. Comparing the detection behavior of the gymnasts group relative to the control group yielded no interesting dissimilarities.

A second concern is whether misalignment of inertia degraded task performance to a significant degree. Figure 5.12(b) plots the group average

performance level as a function of the alignment parameter first for the entire data set and then by number of degrees of freedom. The first observation is that when ignoring effects due to the degrees of freedom, there is little detectable effect of alignment on task performance. This sub-result held between the gymnasts and control group also. Perhaps more surprising is the negative effect that alignment had on three degree of freedom motion. Here a better alignment appears to produce significantly *worse* performance in the data. Examination of the posterior distribution over the precision parameter showed that these differences were significant, though at this point there is neither an *a priori* nor an *a posteriori* hypothesis to explain this effect. Such is the nature of real life data, especially where human subjects are concerned.




Figure 5.12. Alignment detection (a) and performance based on alignment (b) for all subjects.

5.5 Implications for Decision Aids and Other Experiments

The data were gathered in an attempt to characterize the human's abilities

in broad terms when observing objects undergoing natural, unforced rotational motion within an inertial reference frame. The data have resulted in several conclusions that can easily be applied to future experiments and application in this area. These can be summarized as:

Experience Additional experience and possibly training does not appear to be helpful for predicting tumbling motion, that is motion having three degrees of freedom. Training for the experiment was adequate, and no learning over the course of the experiment could be observed.

Observing The process of visually observing the position of an object and estimating its rotation rate parameters to some final level of errors happens quickly in the human, probably within the first 0.5 seconds of observation.

Predicting A two region model of prediction errors is proposed. For rotation rates $\omega_r > .12 \text{ rad/sec}$, the squared error in rad^2 will be proportional to the angle $0.5\alpha_p \text{ rad}$ through which the object rotates during the prediction interval t_p . For rates $\omega_r < .12 \text{ rad/sec}$, the squared error in rad^2 is about equal to the one tenth of the prediction time in seconds, that is $.1t^p$. Thus experiments involving slow bodies should use rotation rates slower than $.12 \text{ rad/sec}$. This model is tentative, given the data.

Screen alignment Motion occurring completely within the plane of the screen appears to be processed the same as motion also occurring in depth.

Mass alignment The human has a slight ability to detect misalignment of inertia parameters and visual geometry, and he is more likely to make these indications when observing more complex motion. Except for possible special behavior for three dimensional motion, task performance suffers predictably when this alignment varies, implying the exact mass distribution or alignment of experimental bodies is not a critical part of the experimental design.

6. Decision Aid Design for Satellite Retrieval

We now apply the results of the previous chapters to a practical problem, retrieving a satellite in space. This problem was selected because it has characteristics which make it a good task for decision aiding. First, the prediction of rotational motion, while hard for a human, is easy for the computer. Second, the decision times involved are moderate, allowing time for the user to consider the information presented by an aid.

6.1 Task Description

6.1.1 Problem Statement

In the future, space travel and activities will be more frequent. As this happens, servicing of satellites in space may become more cost effective than simply replacing malfunction satellites with new ones. Retrieving a satellite for servicing

can be a problem, however, especially if the satellite is nutating or tumbling. We wish to propose a decision aiding system that will help astronauts in this retrieval task.

6.1.2 Previous Approaches

The current thinking on this subject is represented in Rice[46] and Hartley *et. al.*[22]. In this approach, a vehicle such as a Manned Maneuverable Unit (MMU) is first positioned near the target satellite. Then, the MMU circles the target at the same rotation rate and about the same axis as the target. Finally the MMU closes in on the target until the grappling fixture of the target is secured to the MMU. Once this is done the MMU can be used to slow the rotation of the target with its thrusters.

This method has the advantage that thruster plume impingement is not a problem since as the MMU closes on the target, the thrusters must still fire away from the target to maintain the spin around the satellite. This method is increasingly difficult as the rotational motion of the target is more complex. For a tumbling satellite, the target may be considered irretrievable by this method until the motion degenerates into a simpler nutation or spin.

6.1.3 Alternative Approach Using Decision Support System

A decision aid has been designed in which normatively derived state estimates are presented to the human operator. Figure 6.1 shows the elements of this system.

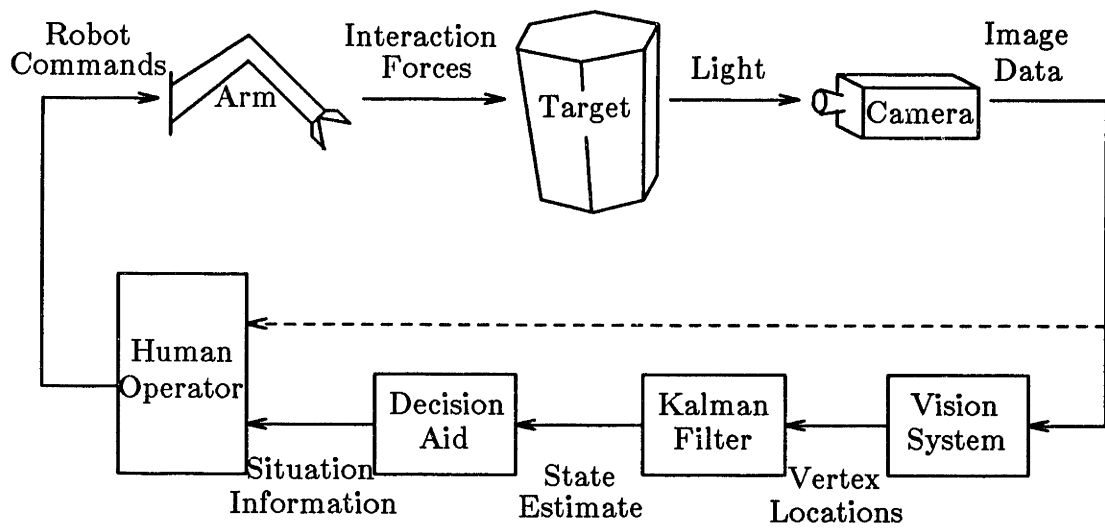


Figure 6.1. The components of a satellite retrieval decision aid based on state estimates.

In our conception of the retrieval mission, the satellite first achieves a parking orbit adjacent to the target satellite. Then a decision support system is used to build up a model of the rotational motion of the satellite in a computer. The operator can work with this model, using it to predict future states of the target or look at past states. Finally, the operator will use a robotic arm to reach out and grapple the satellite as it rotates, using his decision aid to help him select a good opportunity for mission success. Or, the operator gives only the final signal to an automatic arm control and retrieval subsystem.

6.2 Task Analysis

6.2.1 Predictability of Process

Though the rotational motion of the satellite is completely deterministic, tumbling motion can appear to be random and unpredictable to the human

operator. Hence an important way in which the decision aid may be able to help the operator is by providing accurate predictions of future states. Since the state consists of three independent quantities and their first derivatives, communicating a series of future states to the operator may be difficult. Specific modeling uncertainties may reduce the predictability of the state. For example, the exact mass distribution may not be known, or the sloshing action of fuel tanks may make the rigid body motion assumption questionable.

6.2.2 Observability of Process

By observability we mean the degree to which observable data gives unambiguous indications of the process state. The observability of the process state in this system poses no inherent difficulty for the computer, though the algebraic problems of dealing with rotation can be significant. Appendix D gives a method for simulating the dynamics and gives an idea of the computational difficulties. Tanabe, *et.al.* [58] have proposed a system in which positions are automatically estimated by a camera and vision system, as have others. Their system estimates attitude within about 0.5 deg and angular rate error within 0.01 deg/sec.

6.2.3 Decision Time

The target satellites to be captured can be considered to have a nominal rotation rate of 1-3 degrees per second [22]. This means they will completely revolve in two to six minutes. This is a moderately slow system from a decision making point of view since one can consider opportunities for grappling to arise every two to six minutes.

If the rate can be measured to within 0.01 deg/sec, a prediction interval of six minutes will give a thirty six degree interval for position, or about a ten percent error in angular position. Hence the system predictions may span two or three revolutions of possible opportunities before uncertainty in the predictions takes over.

6.2.4 Forgiveness

Another important aspect of this task is that it is unforgiving in the sense of Section 3.2.9. The important final states are divided into success and failure with any of the successful grappling states being equally ranked with the others. Failing to grapple on an attempt can impart additional angular momentum and put the satellite into a worse spin. The forgiveness of the system could be increased by allowing for multiple attempts as part of the task specification. Additional missed attempts will then cause long delays in success, and may damage the target satellite further.

6.2.5 Additional Process Characteristics

This is a single time decision task, since once the satellite is retrieved, there is no need to keep retrieving it. The basic process dynamics are simple without much hierarchy or separable subsystems. We also consider this a target state task because the focus is on the important final grappled state of the satellite. Finally, the task has capturing behavior, since the control method involves observing the process state until the opportunity for grappling exists.

6.3 Design of Decision Aid for Satellite Retrieval

Figure 6.2 shows the simulation and decision aid used to investigate satellite retrieval. The left half of the screen shows the simulated satellite, and in the presented orientation, an "L" shaped grappling fixture is visible. A robotic arm whose end effector matches the "L" shape is also simulated but not shown.




Figure 6.2. Sample display showing simulation and decision aid.

The lower right region of the screen provides the user with a menu of windows of information types to choose from which is described fully in the following sections. The window in the upper right corner is the window of active state information; the horizontal bar displays the time left to make the decision.

6.3.1 On-Line Modeling

The components of the complete decision aiding system are given by Figure 6.1. The elements marked "Vision System" and "Kalman Filter" provide an estimate of the process state together with uncertainty information. Because other researchers are working on the image processing problem [58], we have assumed that such a system is in place and provides noisy measurements of vertex location in terms of two-dimensional viewing screen coordinates.

The Kalman filter was simulated by explicitly implementing one. Unit quaternions are used to represent the angular position of the target while a momentum vector is used to represent the angular velocity. See Appendix B and books such as those by Wertz [67], Junkins [27], and du Val [14] for an introduction to quaternion algebra and its applications. Appendix E gives the formulation for the dynamics and partials necessary to implement the filter.

Inputs to the filter were simulated by computing the screen locations for the present state and adding gaussian noise. By actually implementing the filter, the dynamic behavior of the point estimates and uncertainties was also simulated.

For a linear system, the Kalman filter gains are fixed. If the observations are the deterministic outputs plus gaussian noise, then the point estimate provided by the filter will have noise characteristics that are also gaussian. However, in the nonlinear case where the present estimates are used to compute filter gains, this will not be true. Nor will it be true for the behavior of the covariance matrix representing the uncertainties.

6.3.2 Decision Aid Information Processing - Probabilistic Displays

Figure 6.1 also contains a block marked "Decision Aid" which accepts as input the state estimate data and produces information that is usable by the human operator. The output of the Kalman filter algorithm is a second order approximation to the normative belief state. It contains a vector $\hat{\mathbf{x}}$ representing the point estimate together with an uncertainty represented by the covariance matrix \mathbf{P} . Since the task is unforgiving, the use of probabilistic information is suggested. The additional information requirements however, conflict with the already difficult information requirements using mean value displays.

To assess what information is most useful, the decision aid used in the experiment had several different displays summarizing the decision situation that could be called up by the subject at different times. The displays were organized into two dimensions by rows and columns in the menu: the form of the information and the time scale over which it was displayed. The state estimate information was condensed and displayed to the user in three separate ways that we call pictorial, grappleability, and acceptability displays. Figure 6.3 shows the various displays used in the experiment.

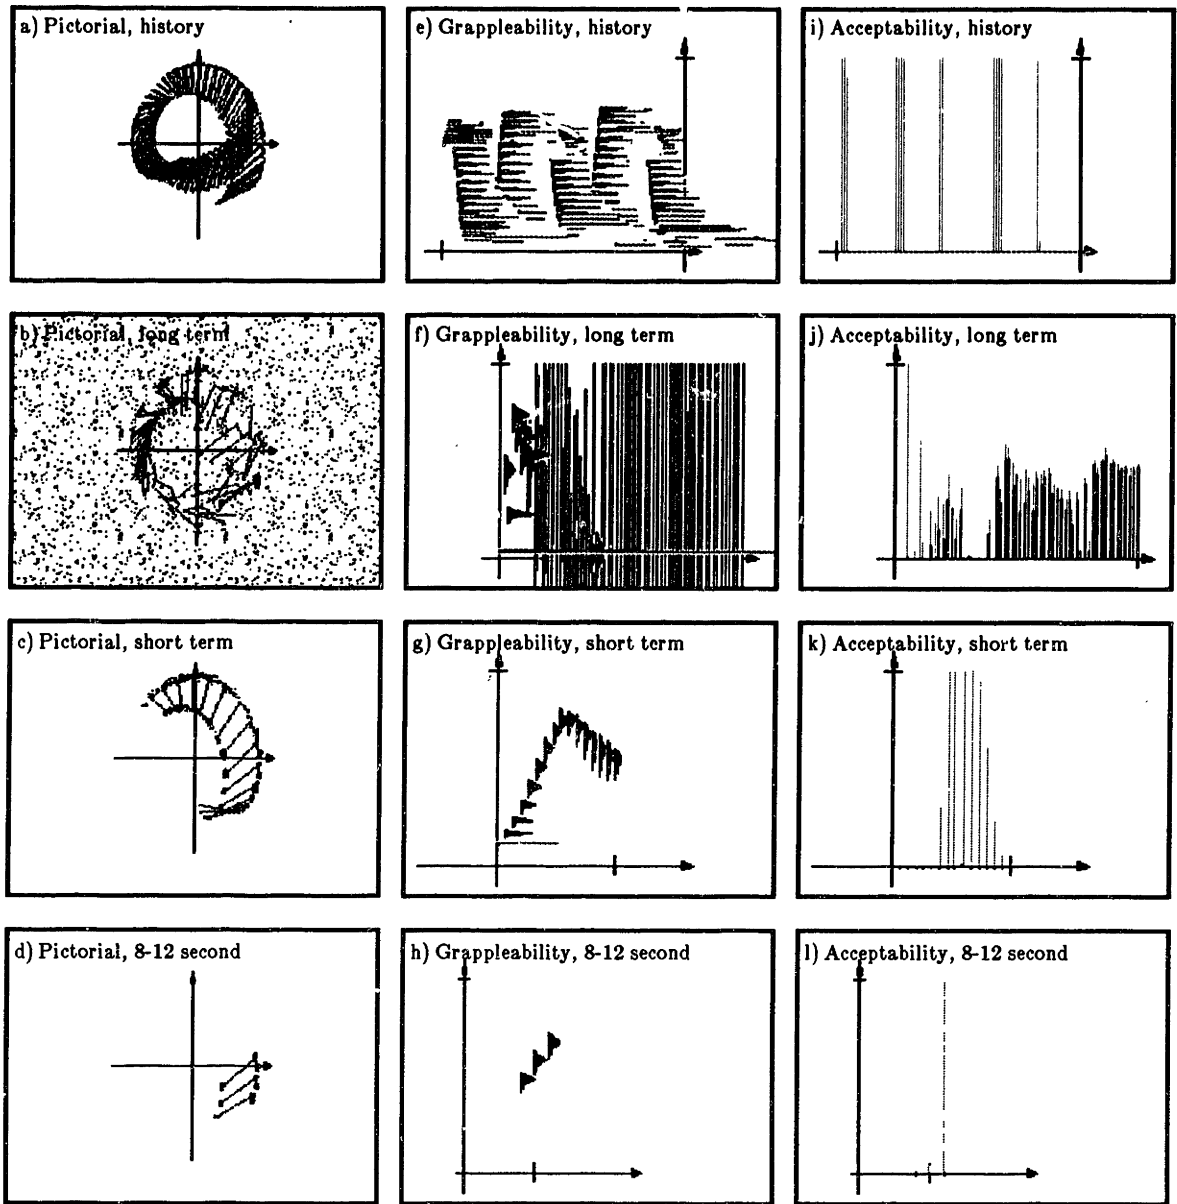


Figure 6.3. Information displays used in the experiment. By column they are pictorial, goodness, and acceptability. By row from top to bottom they are historical, long term prediction, short term prediction, and detailed prediction. The displays are more fully explained in the text.

6.3.3 Pictorial Display - Display Integration

Since the state variables themselves are highly integrated, an integrated display is suggested. In the pictorial displays shown in Figure 6.3(a-d), the estimate and uncertainty are mapped into two dimensional viewing screen coordinates using reference points of the object. If \mathbf{r} is a reference point in the target frame, and $\mathbf{s}(\mathbf{r}, \mathbf{q})$ is an expression for the screen coordinates of \mathbf{r} given position \mathbf{q} , then the screen point \mathbf{x}_s and uncertainty \mathbf{P}_s in terms of a belief state $\hat{\mathbf{q}}$ and uncertainty \mathbf{P} can be computed as

$$\begin{aligned}\mathbf{x}_s &= \mathbf{s}(\mathbf{r}, \hat{\mathbf{q}}) \\ \mathbf{P}_s &= \mathbf{S}\mathbf{P}\mathbf{S}^T\end{aligned}$$

where

$$\mathbf{S} = \left. \frac{\partial \mathbf{s}(\mathbf{r}, \mathbf{q})}{\partial \mathbf{q}} \right|_{\mathbf{q}=\hat{\mathbf{q}}}$$

Appendix E gives relevant formulae for the screen mapping function $\mathbf{s}(\mathbf{v}, \mathbf{q})$ and partials \mathbf{S} .

Practically, there is the problem of displaying the screen uncertainty \mathbf{P}_s , which is a two by two covariance matrix. Points of equivalent uncertainty represented by the covariance matrix will lie along an ellipse, so the approach taken here is to plot an ellipse of equal uncertainty. This may be done by generating points \mathbf{x} which lie along the unit circle, and applying the linear transform,

$$\mathbf{y} = \mathbf{T}\mathbf{x}$$

where

$$\mathbf{T}\mathbf{T}^T = \mathbf{P}_s$$

The matrix \mathbf{T} can be determined by doing a Cholesky decomposition or singular

value decomposition of the matrix \mathbf{P}_s . The transformed \mathbf{y} values will then fall on an ellipse of equal uncertainty and will bound a given fraction of the points. In Figures 6.3(a-d), the uncertainty of reference point locations have been plotted in this fashion as shaded ellipses. As predictions are made farther into the future, uncertainty will increase and these ellipses will grow very large, as they have for long term predictions in Figure 6.3(b).

6.3.4 Grappleability, or Goodness Display - Hierarchical Element

The restructuring of the capture problem into a single time decision task by the addition of computer based observers and automatic arm controlling subsystems has implicitly introduced hierarchical components in the decision aid.

We defined an objective function that can be loosely called "grappleability," or simply "goodness." This function conveys in a single variable the degree to which a certain process state represents an ideal grappling opportunity. The displays of Figure 6.3(e-h) plot grappleability in the vertical dimension as a function of time in the horizontal dimension. This function reduces the information demands on the operator by transforming the decision problem from that of observing and processing six independent variables to that of processing a single process state indication, grappleability. If the function matches some human processing that would be done on the unprocessed information, then computing values for this function would reduce workload without reducing human decision making performance.

The grappleability function is computed as follows. If \mathbf{q} is a unit quaternion representing the current orientation relative to the most desirable orientation, and rotation about the z axis is not important to task success, a suitable scalar valued objective function $g(\mathbf{q})$ is given by,

$$g(\mathbf{q}) \equiv q_1^2 + q_2^2$$

This will measure the square of the sine of the angle through which the body must be rotated to achieve alignment, z axis excluded.

To display beliefs about grappleability, the belief state approximation must be mapped to uncertainty about the grappleability. This may be done by linearizing the objective function about the current estimate. If $\hat{\mathbf{q}}$ and \mathbf{P} represent the belief state, then these may be mapped into a mean \hat{g} and variance v_g in the objective function space by,

$$\begin{aligned}\hat{g} &= g(\hat{\mathbf{q}}) \\ v_g &= \mathbf{G}\mathbf{P}\mathbf{G}^T\end{aligned}$$

where

$$\mathbf{G} \equiv \left. \frac{\partial g(\mathbf{q})}{\partial \mathbf{q}} \right|_{\mathbf{q}=\hat{\mathbf{q}}}$$

The computed mean and variance are then used as parameters of a normal distribution representing beliefs about future values of the grappleability function to generate the plots shown in Figure 6.3(e-h).

6.3.5 Acceptability Display

A third type of display, which we call an acceptability display, was available as shown in Figure 6.3(i-l). These present the probability that the orientation will

be within a certain range of orientations that are acceptable for grappling. Because the boundaries of the acceptability region correspond to physical limitations of the robot arm and are hard constraints, the reward structure is unforgiving and a display of acceptability should be useful in theory.

For this type of case, a mapping may be known or approximated that gives the probability of an event a as a function of the state \mathbf{x} , which is denoted $P(a|\mathbf{x})$. This may be mapped into the probability of event a given both the present or predicted point estimate $\hat{\mathbf{x}}$ and uncertainty \mathbf{P} by,

$$P(a|\hat{\mathbf{x}},\mathbf{P}) = \int_{\mathbf{x}} P(a|\mathbf{x}) (\mathbf{x}-\hat{\mathbf{x}})^T \mathbf{P}^{-1} (\mathbf{x}-\hat{\mathbf{x}}) d\mathbf{x}$$

In our case the function $P(a|\mathbf{x})$ takes on the value 1 for \mathbf{x} which are within the reach of the arm, and 0 for states that are beyond the workspace of the arm. By first transforming the beliefs to a suitable single dimensional subspace, *e.g.* goodness, this formula can be evaluated explicitly and the computed probabilities may be displayed as in Figure 6.3(i-l).

6.3.6 Predictor Displays and Historic Displays

The high degree of predictability in the process dynamics lends itself to the use of predictor displays. For each form of the state information available, predictor displays having three different resolutions were used in the experiment. These were called historic, long term, short term, and 8–12 second displays, with each occupying a row from top to bottom in Figure 6.3. The long term predictions showed estimates of the future spaced at ten second intervals for five minutes and were updated every ten seconds. The short term predictions were estimates

spaced at two seconds for thirty seconds, and were updated every two seconds. The 8–12 second displays showed predictions of the process state at eight, ten, and twelve seconds into the future, and were updated five times per second.

In addition to the predictor display, historical displays of estimates spaced at two seconds for the previous five minutes were available. This type of display gives the operator provides a statistical summary of the past states of the process, so he may judge the degree to which a presented opportunity is unexpectedly good. This increases the effective period of outlook for the operator and should therefore improve his decision making. Historical displays are generated by saving past estimates and will have constant levels of uncertainty over the period if the filter is in steady state.

6.4 Experiments with Decision Aid

Experiments were carried out with a simulator and were designed to answer the following questions.

- Are probabilistic information displays more helpful than mean value displays? Because of the unforgiving task structure and moderately long decision times, we would expect yes.
- What are the various timescales of estimates used by the decision maker.
- What form of information is used by the operators, and how does this affect their decision making performance?
- How do the aid modeling errors affect operator decision performance, as well as his perception and use of the decision aid?

- Is the concept of grappleability a usable concept for the operator? By reducing workload, the and should improve decision performance.
- How helpful is the overall decision aiding system? As a combination of the best components, the system should be helpful. Unanticipated interactions among components or excessive workload may reduce the effectiveness of the synthesis, however.

6.4.1 Experimental Design

One experimental run consisted of six segments of fifteen trials each. There was one training segment which was run first, followed by one segment corresponding to each entry in Figure 6.4. The order of the last five segments was varied among subject to remove order effects.

Condition	Aid?	Aid Modeling Errors	Remarks
No aid	no	-	
Bad-	yes	1×human	Point est. only
Bad+	yes	1×human	Full state est.
Good	yes	.05×human	Full state est.
Perfect	yes	0	

Figure 6.4. Five conditions were tested in the experiment.

Except for the training segment, each segment of fifteen trials was broken into five groups of three trials in which a one DOF, a two DOF, and a three DOF trial were presented in a random order. During a trial, the subject had five minutes during which time he was told he must make a single commit decision when he determined the grappling opportunity was the best it would be during that interval. In all cases a direct view of the target was available. When aid information

was present, he was free to use aid information or direct information.

In a decision aided segment, the first group of three trials was presented with only pictorial information available, the second group of three trials was presented with only grappleability information available, and the third group was presented with only acceptability information available. The last two groups, consisting of six trials, were presented with all types of information available. In the training condition, one trial was presented with no aid, one trial was presented for each information window with only that information available, one trial was presented with all information available, and the final trial was presented with no aid information. When more than one information type was presented, the subject chose from among them, and these choices were recorded for later analysis.

Subjects were seated in front of a computer generated display, a sample of which has been reproduced in Figure 6.1. The subject indicated his irrevocable commit decision by pressing a button on a mouse input device. When a decision was made, a simulation of a robotic end effector was displayed which allowed the user to observe the results of his decision. The arm was servoed to the present estimate of grappling fixture position and velocity in order to "grapple" the simulated target. At the completion of the automatic sequence, the subject received a score representing his decision percentile, *i.e.* that fraction of the decision times for which his decision was superior. He was instructed to behave so as to maximize this percentile score.

Subjects were graduate students in engineering, and three subjects participated in the experiment. In all parts of the experiment, the investigator was available and freely answered questions about the displays and their meanings, but refrained from giving advice about their use or information about the goals of the experiments. The guiding principle in answering the questions was to provide the same level of understanding that would be present in a real system. Subjects were encouraged to discuss their strategies and observations as they performed the experiment, and they were asked for their concept of the investigator's motivation at the completion of the experimental trials.

6.4.2 Initial Conditions and Aid Modeling Errors

In all cases the nominal rotation rate of the target was 0.12 rad/sec. The simulated body was rectangular with dimensions in the ratio 1–2–3 and had principle moments of inertia of 1–.54–.27 along the x , y , and z axes respectively. To produce one DOF motion, the "true" angular momentum vector was initially aligned either along the x -axis or body z -axis in the body frame. To produce two DOF motion, the angular momentum vector was $\pi/8$ radians from the x -axis or z -axis, while for three DOF motion, it was $\pi/8$ radians from the y -axis. The grasping face was on a surface whose normal was along the body y -axis.

The aid modeling errors in Figure 6.4 are reported relative to the human's natural predictive errors. Typical human error growth characteristics were determined in Chapter 5 and are summarized by Figure 5.10. To vary the aid modeling errors, the initial error of the momentum vector estimate was varied, and a small rotational disturbance was applied to the actual inertia tensor. The values

for these disturbances were determined by observing the effects of known errors on simulations until a desired error growth rate characteristic was achieved. For aid modeling errors equal to $1 \times \text{human}$, the momentum vector components were altered by 6%, and a rotational disturbance of .14 radians was applied to the inertia tensor, while for modeling errors equal to $.05 \times \text{human}$, the alterations were 3% and .07 radians.

6.4.3 Performance Measures

The goal of the task was to grapple the target such that the grappling face of the target was as closely aligned as possible with the viewing plane. Thus grappling in a "flat" orientation is preferred to grappling in one that is not "flat." The grappleability function of Section 6.3.4 was also used as the scoring function.

For a given decision situation, the value of the scoring function can be plotted as a function of the possible actions. The actions in this task are uniquely indexed by the time t meaning the scoring function $c(t)$ is a simple function of t . Figure 6.5 shows a possible record of the value of the scoring function and how three performance measures are derived from it.




Figure 6.5. Computation of three performance measures from scoring function and decision data. The decision is made at t_d .

For a given decision situation, the scoring function will take on some set of values over the duration of the decision window $t \in [0, t_w]$. At some point t_g the value of the scoring function will take on a global maximum. If the user had a perfect model of the process, a perfect understanding of the objective function, perfect estimates, and could project the estimates forward using the model without error for at least t_w seconds, then he would know this plot and could select the global maximum as his decision. In practice he will have imperfections and will therefore select some sub-optimal decision at time t_d . We may associate a global loss with his decision $e_g = c(t_g) - c(t_d)$. This is the difference between the utility that he could have gained were he perfect and that which he actually gained.

Similarly, there may be a locally optimal decision time which is shown by t_l in Figure 6.5. This represents the loss associated with choosing an action compared with choosing nearby actions and is computed by $e_l = c(t_l) - c(t_d)$.

Finally, the decision made will be better than some of the decisions available and worse than others. If t_d is the decision time selected by a subject, then

the hatched regions will represent areas when the decision would have been better. This leads to a percentile score, which is the fraction of possible decisions for which the subject's decision was better.

6.4.4 Results

The results of the experiment using the three different performance measures are summarized in Figure 6.6. Regardless of the specific performance measure used, the perfect aid was superior to no aid, and degraded aids were worse than no aid. For global squared error and local squared error, point estimates produced better decision making than displays that used uncertainties. An interesting feature of these plots is that local error was smallest for the only point estimate condition - the bad aid with point estimates. A possible explanation is that the human will only use point estimates in any case, so when these are available directly workload is reduced.




Figure 6.6. Global squared error loss, local squared error loss, and percentile by experimental condition. Plotted are the sample mean and standard error of the observed data.

Another interesting result is the global performance by the motion degrees of freedom. This is shown in Figure 6.7. Here it seen that for simple tasks, that is one DOF tasks, the decision aid is not necessary, and even reduces decision performance slightly. For difficult cases of two and three DOF motion, the decision aid was useful, since the human has no ready algorithms for computing the long term future states of the process.




Figure 6.7. Global squared error as a function of the decision aid and the motion degrees of freedom.

An interesting comment was made by one subject regarding the use of low quality information provided by the decision aid. When asked if he thought it helped him, he said,

"My guess is it did, but I am not sure. Whenever the information was there I just wanted to use it."

This suggests a picture of the displayed information as being an active element that calls out to the operator to be used, rather than the having the passive, inert nature that it is most often given. Further support is found when we recognize that the presence of predictive displays resulted in hasty decisions. The following qualitative observations were also made:

- When point estimates only available, operators still felt comfortable, and in fact their performance improved over the case when uncertainty information was also available. Apparently subjects find it easier to estimate

uncertainty themselves than to accept uncertainty information directly.

- The task is actively reshaped in terms of the available measurements. If the probability of an acceptable range of endpoint angles is displayed, the subject will perform the task in a way that tends to maximize this probability.
- Much of the human's activity in the task was directed towards finding patterns in the environment. Frequently historic displays were used to examine presence or absence of periodicity - a concept that is not explicitly modeled by the state estimation based decision aid.
- Information about past values is less active than information about future values. Subjects were less likely to make hasty decisions when using historic information than when using predictive information.
- The decision aid can have significant emotional effects on the user. For example, one subject reported that the aid gave him a feeling of confidence in making his decisions, and allowed him to manage his own attention resources much better than when it was not available. Another example is the undue sense of urgency that caused hasty actions in several cases. As these ultimately affect the quality of the decisions made, they should not be passed over in real systems.

7. Conclusions

7.1 Summary

The following list provides a brief summary of the major conclusions and contributions of this thesis:

- The general decision aiding problem was examined, and the problem of objectively evaluating decision aids was identified. To deal with this problem, subjective analysis was used in combination with studies of decision making in a controlled laboratory environment.
- The use of normatively derived state estimates as a decision aid was proposed, and experiments were carried out to establish its usefulness for simple systems. An aid that presented state estimates directly was more useful for plants with high output uncertainty.
- A simplification model of the human decision maker was developed in which incoming information is simplified before it modifies the beliefs of the

decision maker. When presented with a distribution, for example, the decision maker will only use a point estimate such as the mode. Simplification can also refer to understanding, *i.e.* when simplified dynamics are assumed in place of true dynamics, or using a rote procedure.

- A characterization of human ability to observe and predict three dimensional rotation was performed. It was found that for a wide range of rotation rates, prediction errors were proportional to subtended angle. Errors were not a function of the observation interval, indicating that for this type of highly visual dynamic observation task, human estimation is complete within 0.5 seconds of presentation, the shortest interval used in the experiment.
- A decision aid for grappling a tumbling satellite was constructed and evaluated. It was determined that as aid modeling errors decrease, decision performance increases. In addition, when aid modeling errors were on the order of the human's, decision performance was worse than the case when no decision aid was present.

7.2 Implications for Research in Related Areas

We now highlight some other areas in which the work presented in this thesis might have implications.

7.2.1 Use of Probability Information in Supervisory Control

Experiments showed that while the use of a distribution over the state space is recommended by the expected utility criterion, human decision makers are

unable or unwilling to use more than point estimates for all but the simplest processes. While nature abhors a vacuum, humans abhor uncertainty. Often these point estimates will be reduced further from estimates over a continuous variable to estimates over a partitioned variable. Though the number of systems to which this statement applies is clearly large, the effect that it should have on the design of man-machine systems is less clear.

It does serve to emphasize the limited degree of understanding which a novice can be expected to possess about a system. We say novice here because in the experiments reported subjects were typically given one to two hours of training with the process before measurements were taken, and this contrasts sharply with the four years of training a power plant operator receives or the 2000 hours of direct flying experience a commercial pilot must have to fly a passenger aircraft. In contrast to novice behavior, the limits to understanding that can be achieved by humans given enough training appear to be few, if they exist at all.

Yet, all other things being equal, we would consider a man-machine system that requires less training to be superior to one that requires more training, so the understanding that a novice brings to a system is important. We may decide that training operators to understand distributions is worthwhile. It seems more appropriate to say that there may be advantages to using distributions in the control systems of the future, and then communicating only the low order statistics of the most important variables to the human, for his immediate understanding.

7.2.2 The Use of a Poor Decision Aid - Active Information

In an experiment with a decision aid constructed for satellite retrieval, the modeling errors of the decision aid were varied. Results showed that the presence of the poor decision aid produced worse performance than having no aid.

The primary mechanism behind this seems to be that the human begins to feel that he must use direct information in addition to the decision aid information, and so he tries to share his attention between the two sources. Instead of the aid augmenting his own beliefs, the aid then interferes with the task as he would otherwise do it.

In addition to causing problems due to switching between sources, the presence of the poor decision aid was over-used by the subjects. As one subject put it, when the information was there he felt he just had to use it. In addition, depending on what was being displayed, the subjects had emotional reactions of pressure, comfort, control. The activity of the information as well as its information value must therefore be taken into account when deciding to display it. We simply cannot continue to pile on information to our operators and require more training - eventually it will require a life's worth of training to control some of our systems, and the operator will die just before he gets any experience controlling the real process.

7.2.3 The Evaluating Problem of Multi-Attribute Decision Making

The objective evaluation problem is inherent in the decision aiding domain.

In this effort we have worked with single attribute decision making problems, but multi-attribute decision problems have the same characteristic. Typically one will first face this problem when sitting down to provide a decision aid for a real problem. At some point he will ask, "How am I going to show that it works?" He will then see that if he is able to assign a scoring function, the aid is not necessary because the decision problem can be solved explicitly. He must then resort to the subjective evaluation method that we have used. Here decision aiding *techniques* were first evaluated in the laboratory using a known, simulated process. When their usefulness was established in the laboratory, they were applied to the real system for which the model is unknown.

In our view, the uncertainty of a model in single attribute decision problem and the multiple incommensurable attributes of the multi objective problems arise from the same cause: *modeling convenience*. The uncertainty arises because, even though the problem may embody a single deterministic mechanism that produces a single behavior, it is too complicated to identify each element. Uncertainty is then used as a convenient and orderly way of considering all possible outcomes simultaneously and to varying degrees. In the multi attribute case, the decision problem is single attribute in nature, but we find it extraordinarily difficult to determine all the interactions that lead to various values of the fundamental utility. A furniture company is interested only in profit, yet all the interactions that relate the color of their product to profit are too difficult to model explicitly, so the decision is made in its decomposed form with the human providing his own approximate functional mappings to arrive at a decision that is

near optimal with respect to the single objective, profit.

To evaluate a decision aid in a multi-attribute environment, the same subjective evaluation methods must be used. A known model is constructed with a single attribute utility criterion. The decision aid and decision maker will not have direct access to the model or the criterion, but the information available to the human must be such that if he were a perfect observer and modeler, he could infer the model, the state, and the utility function. The techniques are then evaluated with respect to this known criterion. Once the technique is demonstrated in the laboratory, it is applied to real decision problems of a similar nature on faith.

7.3 Recommendations for Additional Research

7.3.1 Additional Application Areas and Poorly Modeled Systems

Though the concept of a decision aid based on normatively derived state estimates was applied to the problem of capturing a rotating satellite, more applications can be examined. Some possible application areas for which models can usually be constructed are power plants, ships, airplanes, and space. Some less easily modeled systems include economic systems, political systems, and military systems. Examples of a less serious nature include crossing a street, choosing a soulmate, or controlling one's own study habits.

Since we can expect to have systems at the limits of our own modeling ability for some time to come, and we will want human operators present to respond

to unanticipated circumstances, it is important to develop techniques that are useful in these systems, *i.e.* we must expect the unexpected. Research may be done about the common features of unexpected events, such as accidents or disruptions, how people should respond, and how they do respond to them. The subjective evaluation method of using known processes and models degraded in known ways may lead to generic aids for unexpected events.

7.3.2 Probabilities, Training, and Training in Probabilities

Since novices do not adapt well to the presentation of probabilistic state information, training in interpreting probabilities for decision making may be important. Though this goes against earlier statements that more training is not always the answer, it underlines the need for more understanding of what role probability information should play in future control systems.

As we have said, the abilities of trained operators differ considerably from those of untrained operators, and it may be useful to understanding how a well trained operator becomes well trained. Research could consist of taking an arbitrary task and watching a subject over an enormously long period, and doing some characterization. By knowing the characteristics of operators more fully, we may be able design the man-machine interface and training as an integrated whole.

7.3.3 Simplifications in Theory and Practice

It remains unclear how the human decision maker selects a decision algorithm that simplifies data yet retains the essential characteristics to achieve

moderately good decision making behavior. Yet this is an important part of intelligent activity. To understand how this is done would, in essence, be to understand the modeling process. An engineering model presents a collection of ideas which contain the essential behavior of the thing being modeled, and contain little else. To understand simplifications, then, is to understand the process of model building and model using, which are essential to intelligent behavior.

Making simplifications allows the human to deal with complex systems using limited resources. We are rapidly producing computers with ever-expanding lists of available computing resources. Yet we have devoted little effort to seeing how the resources of an intelligent machine can be managed - tasks simplified, problems reduced, patterns found and exploited. Combining the expanding artificial resources with human-like capabilities of resource management could produce systems of enormous power, use and reliability.

Appendix A. Vectors and Rotation Matrices

We quickly review some aspects of the vector and matrix formulation used throughout the work. To represent the angular relationship between two reference frames, a direction cosine matrix or rotation matrix is used. This is a 3×3 proper real orthogonal matrix that provides a convenient way to transform locations expressed in one reference frame to another coordinate frame. A computationally useful property of the rotation matrix \mathbf{C} also expresses the constraints imposed for \mathbf{C} to be a rotation matrix:

$$\mathbf{C}\mathbf{C}^T = \mathbf{1}$$

or equivalently,

$$\mathbf{C}^{-1} = \mathbf{C}^T$$

Additionally, $\det(\mathbf{C})$ must equal 1 for \mathbf{C} to express a right handed coordinate system. If \mathbf{C} is a rotation matrix and \mathbf{x} is a vector expressed in the global coordinate frame, then

$$\mathbf{y} = \mathbf{C}\mathbf{x}$$

is the same point expressed in the body coordinate system. Of course given \mathbf{y} expressed in body coordinates, one can determine \mathbf{x} from

$$\mathbf{x} = \mathbf{C}^T\mathbf{y}.$$

Two useful ways of thinking of the matrix \mathbf{C} are in terms of column vectors and row vectors:

$$\mathbf{C} = \begin{bmatrix} \mathbf{n}_x & \mathbf{n}_y & \mathbf{n}_z \end{bmatrix} = \begin{bmatrix} \mathbf{u}_x^T \\ \mathbf{u}_y^T \\ \mathbf{u}_z^T \end{bmatrix}.$$

The vector \mathbf{n}_x expresses the unit vector along the x axis of the global frame in terms of the body coordinate system. Similarly, the vector \mathbf{u}_x expresses the unit vector along the x axis of the body coordinate frame in terms of global coordinates.

In this view the \mathbf{n}_i and \mathbf{u}_i are unit vectors also satisfying constraints

$$\begin{aligned}\mathbf{n}_x^T \mathbf{n}_y &= \mathbf{n}_y^T \mathbf{n}_z = \mathbf{n}_z^T \mathbf{n}_x = 0 \\ \mathbf{u}_x^T \mathbf{u}_y &= \mathbf{u}_y^T \mathbf{u}_z = \mathbf{u}_z^T \mathbf{u}_x = 0\end{aligned}$$

A rotation matrix may be used to represent the time varying angular position of a body. The concatenation of two rotation matrices will also be a rotation matrix whose reference frame has an angular displacement representing successive displacements of the two matrices. If \mathbf{C} represents the present angular position of a body and \mathbf{D} represents a small x axis rotation, the new position \mathbf{C}' resulting when the body is rotated about its own x axis is given by

$$\mathbf{C}' = \mathbf{D}\mathbf{C}$$

In a complementary way, the new position \mathbf{C}' resulting when the body is rotated about the global x axis is given by

$$\mathbf{C}' = \mathbf{C}\mathbf{D}$$

Since rotation requires only three independent quantities to specify, the rotation matrix contains six redundant quantities that are determined from the constraints outlined above. After many computations involving rotation matrices, it is possible that these constraints are no longer satisfied. Here the matrix \mathbf{C} can be renormalized to \mathbf{C}' . If \mathbf{C} is almost orthonormal already, the following will very nearly renormalize \mathbf{C} .

$$\begin{aligned}\mathbf{n}'_x &= \frac{1}{|\mathbf{n}_x|} \mathbf{n}_x - \frac{1}{2} \left((\mathbf{n}_x^T \mathbf{n}_y) \mathbf{n}_y + (\mathbf{n}_x^T \mathbf{n}_z) \mathbf{n}_z \right) \\ \mathbf{n}'_y &= \frac{1}{|\mathbf{n}_y|} \mathbf{n}_y - \frac{1}{2} \left((\mathbf{n}_y^T \mathbf{n}_z) \mathbf{n}_z + (\mathbf{n}_y^T \mathbf{n}_x) \mathbf{n}_x \right) \\ \mathbf{n}'_z &= \frac{1}{|\mathbf{n}_z|} \mathbf{n}_z - \frac{1}{2} \left((\mathbf{n}_z^T \mathbf{n}_x) \mathbf{n}_x + (\mathbf{n}_z^T \mathbf{n}_y) \mathbf{n}_y \right)\end{aligned}$$

Appendix B. Quaternions as an Angular Position State Variable

Because the rotation matrix has nine variables to represent three independent quantities, its use can be cumbersome. In contrast, three variable parameterizations such as Euler angles or Gibb's vectors have no redundant variables but all pose difficulties with singularities for some positions. A representation that avoids some of these difficulties is the quaternion representation. Some properties of quaternions are reviewed here. A quaternion q has four components

$$q = \begin{bmatrix} q_1 & q_2 & q_3 & q_4 \end{bmatrix}^T$$

The quaternion q may also be treated as a composition of a real scalar q_4 and a so-called hyperimaginary 3-vector \mathbf{q} .

$$q = (q_4, \mathbf{q})$$

This allows for convenient expression of quaternion operations in terms of well known scalar and 3-vector operations. For example, the operations of quaternion conjugate q^* , negation $-q$, norm $|q|$, sum $q+q'$, and product qq' are

$$q^* = (q_4, -\mathbf{q})$$

$$-q = (-q_4, -\mathbf{q})$$

$$|q| = q_4^2 + \mathbf{q} \cdot \mathbf{q}$$

$$q+q' = (q_4+q'_4, \mathbf{q}+\mathbf{q}')$$

$$qq' = (q_4q'_4 - \mathbf{q} \cdot \mathbf{q}', q_4\mathbf{q}' + q'_4\mathbf{q} + \mathbf{q} \times \mathbf{q}')$$

Given quaternions q , r , and s , the following properties of quaternion operations are useful

$$qq^* = (|q|, 0)$$

$$q^{-1} = \frac{1}{|q|} q^*$$

$$(qr)^* = r^*q^*$$

$$(qr)s = q(rs)$$

A scalar a may be represented by quaternion $\mathbf{a}=(a, \mathbf{0})$, and a vector \mathbf{v} may be represented by quaternion $\mathbf{v}=(0, \mathbf{v})$. This leads to the following definitions

$$aq \equiv aq = (aq_4, a\mathbf{q})$$

$$q\mathbf{v} \equiv q\mathbf{v} = (-\mathbf{q}\cdot\mathbf{v}, q_4\mathbf{v} + \mathbf{q}\times\mathbf{v})$$

$$\mathbf{v}q \equiv \mathbf{v}q = (-\mathbf{v}\cdot\mathbf{q}, q_4\mathbf{v} + \mathbf{v}\times\mathbf{q})$$

Unit quaternions are used to represent rotations. For every unit quaternion q , there can be found an associated unit vector \mathbf{u} and angle θ such that

$$q = (\cos(\theta), \sin(\theta)\mathbf{u}).$$

Unit quaternions will then satisfy

$$|q| = 1$$

$$q^{-1} = q^*$$

Note that the requirement that q be a unit quaternion provides a single constraint to offset the single redundant parameter. If the unit quaternion q represents the angular displacement of a rigid body, and \mathbf{w} is a vector in global coordinates, the corresponding point \mathbf{v} expressed in body coordinates can be computed using quaternions $\mathbf{w}=(0, \mathbf{w})$ and $\mathbf{v}=(0, \mathbf{v})$ as

$$\mathbf{v} = q\mathbf{w}q^*$$

Similarly, \mathbf{v} expressed in body coordinates may be resolved to world coordinates \mathbf{w} by the quaternion product

$$\mathbf{w} = q^*\mathbf{v}q$$

Appendix C. Commonly Used Matrix Functions

In deriving equations of motion and estimation equations, several matrices appear frequently. Using these matrix functions simplifies later notation and often simplifies analysis because of some useful properties, which we also present. First, when working with vectors and direction cosine matrices, the following definition is useful:

$$\mathbf{\Omega}(\boldsymbol{\omega}) \equiv \begin{bmatrix} 0 & \omega_z & -\omega_y \\ -\omega_z & 0 & \omega_x \\ \omega_y & -\omega_x & 0 \end{bmatrix} +$$

A few properties of this matrix are useful when taking derivatives of expressions and can result in computationally simple results:

$$\mathbf{\Omega}(-\boldsymbol{\omega}) = -\mathbf{\Omega}(\boldsymbol{\omega}) = \mathbf{\Omega}^T(\boldsymbol{\omega})$$

$$\mathbf{\Omega}(\boldsymbol{\omega})\mathbf{h} = -\mathbf{\Omega}(\mathbf{h})\boldsymbol{\omega}$$

$$\frac{\partial(\mathbf{\Omega}(\boldsymbol{\omega})\mathbf{h})}{\partial\boldsymbol{\omega}} = \mathbf{\Omega}(\mathbf{h})$$

Another matrix used in rotation is the expansion for a finite rotation about a fixed axis. If a reference frame is rotated θ radians about a constant axis whose unit vector is given by \mathbf{u} , then the cosine matrix that expresses this transformation is given by,

$$\mathbf{S}(\theta, \mathbf{u}) \equiv \cos(\theta)\mathbf{1} + \sin(\theta) \begin{bmatrix} 0 & u_z & -u_y \\ -u_z & 0 & u_x \\ u_y & -u_x & 0 \end{bmatrix} + (1-\cos(\theta)) \begin{bmatrix} u_x^2 & u_x u_y & u_x u_z \\ u_x u_y & u_y^2 & u_y u_z \\ u_x u_z & u_y u_z & u_z^2 \end{bmatrix}$$

Some expressions involving this matrix are,

$$\mathbf{S}(\theta, \mathbf{u}) = \cos(\theta)\mathbf{1} + \sin(\theta)\mathbf{\Omega}(\mathbf{u}) + (1-\cos(\theta))\mathbf{u}\mathbf{u}^T$$

$$\mathbf{S}(-\theta, -\mathbf{u}) = \mathbf{S}(\theta, \mathbf{u})$$

$$\lim_{\theta \rightarrow 0} \mathbf{S}(\theta, \mathbf{u}) = \mathbf{\Omega}(\theta\mathbf{u})$$

When using quaternions, four dimensional matrices can be used to express the operation of a left screw and a right screw. We therefore define the following four by four matrices:

$$\mathbf{Q}^R(\mathbf{q}) \equiv \begin{bmatrix} q_4 & -q_3 & q_2 & q_1 \\ q_3 & q_4 & -q_1 & q_2 \\ -q_2 & q_1 & q_4 & q_3 \\ -q_1 & -q_2 & -q_3 & q_4 \end{bmatrix} \quad \mathbf{Q}^L(\mathbf{q}) \equiv \begin{bmatrix} q_4 & q_3 & -q_2 & q_1 \\ -q_3 & q_4 & q_1 & q_2 \\ q_2 & -q_1 & q_4 & q_3 \\ -q_1 & -q_2 & -q_3 & q_4 \end{bmatrix}$$

These definitions allow quaternion multiplication to be expressed in matrix form and simplify analysis by virtue of the following relations:

$$\begin{aligned} \mathbf{q}\mathbf{q}' &= \mathbf{Q}^R(\mathbf{q})\mathbf{q}' = \mathbf{Q}^L(\mathbf{q}')\mathbf{q} \\ \mathbf{Q}^R(\mathbf{q}^*) &= \mathbf{Q}^R(\mathbf{q})^T \quad ; \quad \mathbf{Q}^L(\mathbf{q}^*) = \mathbf{Q}^L(\mathbf{q})^T \\ \mathbf{Q}^R(\mathbf{Q}^R(\mathbf{q})\mathbf{q}') &= \mathbf{Q}^R(\mathbf{q})\mathbf{Q}^R(\mathbf{q}') \quad ; \quad \mathbf{Q}^L(\mathbf{Q}^L(\mathbf{q})\mathbf{q}') = \mathbf{Q}^L(\mathbf{q})\mathbf{Q}^L(\mathbf{q}') \\ \mathbf{q}\mathbf{q}'\mathbf{q}^{-1} &= \mathbf{Q}^R(\mathbf{q})\mathbf{Q}^L(\mathbf{q}^{-1})\mathbf{q}' = \mathbf{Q}^L(\mathbf{q}^{-1})\mathbf{Q}^R(\mathbf{q})\mathbf{q}' \\ \frac{\partial(\mathbf{q}\mathbf{q}')}{\partial\mathbf{q}} &= \mathbf{Q}^L(\mathbf{q}') \quad ; \quad \frac{\partial(\mathbf{q}'\mathbf{q})}{\partial\mathbf{q}} = \mathbf{Q}^R(\mathbf{q}') \\ \frac{\partial(\mathbf{q}^*\mathbf{q}')}{\partial\mathbf{q}} &= \mathbf{Q}^L(\mathbf{q}'^*) \quad ; \quad \frac{\partial(\mathbf{q}'\mathbf{q}^*)}{\partial\mathbf{q}} = \mathbf{Q}^R(\mathbf{q}'^*) \end{aligned}$$

If \mathbf{q} is a unit quaternion, it can be decomposed into a vector part and a scalar part, $\mathbf{q}=(s\mathbf{u},c)$ such that \mathbf{u} is a unit vector and $s^2+c^2=1$. Here \mathbf{u} represents the Eulerian axis about which the frame is rotated, and $\theta:s=\sin(\theta),c=\cos(\theta)$ is one half the Euler angle of the rotation. A direction cosine matrix can be determined from the quaternion as follows:

$$\mathbf{A}(\mathbf{q}) = \begin{bmatrix} q_1^2 - q_2^2 - q_3^2 + q_4^2 & 2(q_1q_2 + q_3q_4) & 2(q_1q_3 - q_2q_4) \\ 2(q_1q_2 - q_3q_4) & -q_1^2 + q_2^2 - q_3^2 + q_4^2 & 2(q_2q_3 + q_1q_4) \\ 2(q_1q_3 + q_2q_4) & 2(q_2q_3 - q_1q_4) & -q_1^2 - q_2^2 + q_3^2 + q_4^2 \end{bmatrix}$$

In the following it is assumed that $\mathbf{q}=(c,s\mathbf{u})$ is a unit quaternion, and \mathbf{h} is a quaternion representation of a vector \mathbf{h} , i.e. $\mathbf{h}=(0,\mathbf{h})$.

$$\begin{aligned} \mathbf{A}(\mathbf{q}) &= \mathbf{S}(\tan^{-1}(s/c),\mathbf{u}) \\ \mathbf{q}\mathbf{h}\mathbf{q}^{-1} &= \mathbf{q}\mathbf{h}\mathbf{q}^* = \mathbf{Q}^R(\mathbf{q})\mathbf{Q}^L(\mathbf{q}^*)\mathbf{h} = \begin{bmatrix} \mathbf{A}(\mathbf{q})\mathbf{h} & \mathbf{0} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \\ \frac{\partial(\mathbf{q}\mathbf{h}\mathbf{q}^*)}{\partial\mathbf{q}} &= 2\mathbf{Q}^L(\mathbf{h}\mathbf{q}^*) = 2\mathbf{Q}^L(\mathbf{q}^*)\Omega^{\mathbf{Q}}(\mathbf{h}) \end{aligned}$$

where

$$\Omega^Q(\mathbf{h}) = \begin{bmatrix} 0 & h_z & -h_y & h_x \\ -h_z & 0 & h_x & h_y \\ h_y & -h_x & 0 & h_z \\ -h_x & -h_y & -h_y & 0 \end{bmatrix} = \mathbf{Q}^L(\mathbf{h})$$

A useful relationship involving $\Omega^Q(\omega)$ is to note that

$$\frac{\partial \Omega^Q(\omega(x))}{\partial x} = \Omega(\partial \omega(x) / \partial x)$$

Appendix D. Numerical Integration of Equations of Motion for a Rotating Rigid Body

To ensure angular momentum is conserved, the angular momentum is used as part of the state vector. A second order Euler integration method is used and combined with finite rotation expressions to achieve high precision. Then the result is adjusted so that the kinetic energy remains constant before and after the integration. Assume the following values are known:

- \mathbf{h} - the momentum vector at time t
- \mathbf{C} - the cosine matrix at time t
- \mathbf{I}_b - the inertia tensor in the body frame
- t_d - the integration step size.

First, the angular momentum relative to the body frame \mathbf{h}_b and the instantaneous angular velocity in the body frame $\boldsymbol{\omega}_b$ are determined:

$$\mathbf{h}_b = \mathbf{C}^T \mathbf{h}$$

$$\boldsymbol{\omega}_b = \mathbf{I}_b^{-1} \mathbf{h}_b$$

Next, an average angular velocity over the duration of the timestep $\bar{\boldsymbol{\omega}}$ is determined by first determining the rate of change of angular velocity $\dot{\boldsymbol{\omega}}_b$:

$$\dot{\boldsymbol{\omega}}_b = \mathbf{I}_b^{-1} \boldsymbol{\Omega}(\boldsymbol{\omega}_b) \mathbf{C}^T \mathbf{h}$$

$$\bar{\boldsymbol{\omega}} = \boldsymbol{\omega}_b + \frac{1}{2} \dot{\boldsymbol{\omega}}_b t_d$$

Then a finite rotation is determined about a fixed axis and applied to the present cosine matrix to determine the next cosine matrix.

$$\bar{\theta} = |\bar{\boldsymbol{\omega}}| t_d$$

$$\mathbf{u}_\omega = \bar{\boldsymbol{\omega}} / |\bar{\boldsymbol{\omega}}|$$

$$\mathbf{C}^{(-)} = \mathbf{S}(\bar{\theta}, \mathbf{u}_\omega) \mathbf{C}$$

where $\mathbf{S}(\theta, \mathbf{u})$ is the direction cosine matrix which rotates a frame θ radians about an axis \mathbf{u} and is defined in Appendix C.

To achieve long term stability in the numerical integration, the kinetic energy terms must remain constant. An easy way to ensure this is to apply a rotational correction to the cosine matrix to bring the kinetic energy back to the previous value. Therefore, we begin by measuring the change in kinetic energy

T_e due to the integration:

$$\begin{aligned}\mathbf{h}_b^{(-)} &= \mathbf{C}^{(-)}\mathbf{h} \\ \boldsymbol{\omega}_1^{(-)} &= \mathbf{I}_b^{-1}\mathbf{h}_b \\ T_e &= \boldsymbol{\omega}_b^{(-)T}\mathbf{h}_b^{(-)} - \boldsymbol{\omega}_b^T\mathbf{h}_b\end{aligned}$$

Next, the rate of change of the kinetic energy terms for a given small rotational correction ∇_θ is determined and the value of this correction is determined and applied to give the corrected integration $\mathbf{C}^{(+)}$:

$$\begin{aligned}\mathbf{T}' &\equiv \frac{\partial T}{\partial \nabla_\theta} = -2\boldsymbol{\Omega}(\mathbf{h}_b^{(-)})\boldsymbol{\omega}_b \\ \mathbf{u}_\nabla &= \mathbf{T}'/|\mathbf{T}'| \\ \theta_\nabla &= T_e/|\mathbf{T}'| \\ \mathbf{C}^{(+)} &= \mathbf{S}(\theta_\nabla, \mathbf{u}_\nabla)\mathbf{C}^{(-)}\end{aligned}$$

Appendix E. Estimator System Equations Definition

To derive a state estimation algorithm using the extended Kalman filter algorithm, the following system state equations were used. See Gelb [20] for a review of the extended Kalman filter equations. The state variables for angular position were given by a unit quaternion \mathbf{q} while the angular velocity is represented by the angular momentum vector \mathbf{h} . Given values for the present state, the system differential equation is then given by,

$$\mathbf{f}(\mathbf{h}, \mathbf{q}) = \begin{bmatrix} \dot{\mathbf{h}} \\ \dot{\mathbf{q}} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \frac{1}{2}\Omega^Q(\omega)\mathbf{q} \end{bmatrix}$$

where

$$\omega = \mathbf{I}_b^{-1} \mathbf{A}(\mathbf{q})\mathbf{h}$$

Here \mathbf{I}_b is the inertia tensor expressed in the body coordinate frame, and $\Omega^Q(\omega)$ and $\mathbf{A}(\mathbf{q})$ are matrix functions as given in Appendix C. To apply the extended Kalman filter, the system function is linearized about the present estimate as follows:

$$\mathbf{F}(\mathbf{h}, \mathbf{q}) = \begin{bmatrix} \partial\dot{\mathbf{h}}/\partial\mathbf{h} & \partial\dot{\mathbf{h}}/\partial\mathbf{q} \\ \partial\dot{\mathbf{q}}/\partial\mathbf{h} & \partial\dot{\mathbf{q}}/\partial\mathbf{q} \end{bmatrix}$$

$$\frac{\partial\dot{\mathbf{h}}}{\partial\mathbf{h}} = \frac{\partial\dot{\mathbf{h}}}{\partial\mathbf{q}} = \mathbf{0}$$

$$\frac{\partial\dot{\mathbf{q}}}{\partial h_i} = \frac{1}{2}\Omega^Q\left(\frac{\partial\omega}{\partial h_i}\right)\mathbf{q}$$

$$\frac{\partial\dot{\mathbf{q}}}{\partial q_i} = \frac{1}{2}\Omega^Q\left(\frac{\partial\omega}{\partial q_i}\right)\mathbf{q} + \frac{1}{2}\Omega^Q(\omega)\frac{\partial\mathbf{q}}{\partial q_i}$$

where

$$\frac{\partial\omega}{\partial\mathbf{h}} = \begin{bmatrix} \frac{\partial\omega}{\partial h_x} & \frac{\partial\omega}{\partial h_y} & \frac{\partial\omega}{\partial h_z} \end{bmatrix} = \mathbf{I}_b^{-1} \mathbf{A}(\mathbf{q})$$

$$\frac{\partial\omega}{\partial\mathbf{q}} = \begin{bmatrix} \frac{\partial\omega}{\partial q_1} & \frac{\partial\omega}{\partial q_2} & \frac{\partial\omega}{\partial q_3} & \frac{\partial\omega}{\partial q_4} \end{bmatrix} = -2\mathbf{I}_b^{-1} \mathbf{Q}^R(\mathbf{q}^*)\Omega^Q(\mathbf{h})$$

and the matrix $\mathbf{Q}^R(\mathbf{q})$ is as given in Appendix C. In this formulation, if a 3×3 matrix \mathbf{A} is multiplied by a 4×4 matrix \mathbf{B} the former is expanded by setting

$a_{i4}=a_{4i}=0; i=1,3$ and $a_{44}=1$.

Additionally, the system output functions must be specified, and this is where the constraint imposed on the redundant state variable must be explicitly taken into account. The output consists of a set of screen coordinates for each of the m measured vertex locations and an expression for the single constraint. Thus the output function $\mathbf{g}(\mathbf{h}, q)$ is given by,

$$\mathbf{g}(\mathbf{h}, q) = \begin{bmatrix} \mathbf{s}(\mathbf{w}_1) \\ \vdots \\ \mathbf{s}(\mathbf{w}_m) \\ |q| - 1 \end{bmatrix}$$

where the world coordinates of each point are given by

$$\mathbf{w}_i = \mathbf{A}(q)\mathbf{v}_i; \quad i=1, m$$

and these are mapped to screen coordinates by the function

$$\mathbf{s}(\mathbf{w}) = \begin{bmatrix} \cot(\gamma_x)w_x/w_z \\ \cot(\gamma_y)w_y/w_z \end{bmatrix}$$

Here \mathbf{v}_i represents the location of the i^{th} vertex expressed in body coordinates, \mathbf{w}_i is the i^{th} point in world coordinates, and the function $\mathbf{s}(\mathbf{w})$ maps a point \mathbf{w} in world coordinates to a point $\mathbf{s}(\mathbf{w})$ in screen coordinates. As in the system function, the output function must also be linearized about the current estimate:

$$\mathbf{G}(\mathbf{h}, q) = \begin{bmatrix} \mathbf{0} & \mathbf{S}(\mathbf{w}_1)\partial\mathbf{w}_1/\partial q \\ \vdots & \vdots \\ \mathbf{0} & \mathbf{S}(\mathbf{w}_m)\partial\mathbf{w}_m/\partial q \\ \mathbf{0} & 2q \end{bmatrix}$$

where partials are taken for each set of world coordinates

$$\partial\mathbf{w}_i/\partial q = 2\mathbf{Q}^L(q^*)\Omega^Q(\mathbf{v}_i)$$

and the screen mapping function partials are evaluated as

$$\mathbf{S}(\mathbf{w}) = \begin{bmatrix} \cot(\gamma_x)/w_z & 0 & -\cot(\gamma_x)w_x/w_z^2 \\ 0 & \cot(\gamma_y)/w_z & -\cot(\gamma_y)w_y/w_z^2 \end{bmatrix}$$

Appendix F. Bayesian Estimation of Model Parameters

To analyze experimental data, a method proposed and developed by Max Mendel was used. An assumption in this method, is that the data are exchangeable, *i.e.*, their order has no effect. Bayesian analysis is used to determine a posterior distribution over the next data value, x_0 , given the experimental data observed. A posterior distribution over the model parameter(s) is determined only as a side effect of the primary goal of predicting the next observation from previous ones.

We present here a sketch of the computational methods used. The reader is referred to Mendel's work for details [39]. We begin with n data points x_1 through x_n and we denote them as $\mathbf{x}_n = x_1, \dots, x_n$. We assume these data points to be squared norms of a vector having d elements generated from the same normal distribution and having zero mean and unknown precision ω . Selecting precision rather than its inverse variance is one of several decisions that may appear arbitrary but simplify matters greatly by allowing the use of conjugate distributions when updating estimates given new data. In our application $d=3$. A distribution is assumed over the ω representing our beliefs about its value. By assuming ω to be γ_2 distributed, it can be shown that our resulting belief about the x_i is an F distribution.

The posterior distribution over the next data point x_0 given the data \mathbf{x}_n can be shown to be,

$$f(x_0 | \mathbf{x}_n) = f_F(x_0 | 1, d\nu^2, \nu'^{-1}) \quad (\text{F.1})$$

while the posterior distribution over the precision parameter ω

$$f(\omega | \mathbf{x}_n) = f_{\gamma_2}(\omega | \nu', d\nu') \quad (\text{F.2})$$

where

$$\nu' = \frac{\nu v + n\bar{m}(\mathbf{x}_n)/d}{\nu + n},$$

$$\nu' = \nu + n.$$

where n is the number of data points taken, and $\bar{v}(\mathbf{x}_n)$ ν and v are prior distribution parameters whose selection is discussed in the following section, and $\bar{m}(\mathbf{x}_n)$ is the sample mean of the data set \mathbf{x}_n defined in the usual way as,

$$\bar{m}(\mathbf{x}_n) \equiv \sum_{i=1}^n x_i/n$$

The mean $E(x_0)$ and variance $V(x_0)$ of the posterior distribution over the next data point x_0 are determined from the F-distribution in equation (F.1) as,

$$E(x_0) = \frac{v' d\nu'}{(d\nu' - 2)} \quad (\text{F.3})$$

$$V(x_0) = \frac{2v'^2(d\nu')^2(d\nu' - 1)}{(d\nu' - 2)^2(d\nu' - 4)}$$

while the mean $E_{\gamma_2}(\omega)$ and variance $V_{\gamma_2}(\omega)$ of the posterior distribution over the precision parameter ω is determined from the gamma-2 distribution in equation (F.2) as,

$$E_{\gamma_2}(\omega) = \frac{1}{v'}$$

$$V_{\gamma_2}(\omega) = \frac{2}{v'^2 d\nu'}$$

F.1 Selection of Prior Distribution

To use the method, a prior distribution over the precision parameter ω must be specified using the parameters ν and v . The parameter ν is interpreted as an equivalent number of data points to be associated with the prior distribution and is normally set to 1. The parameter v is a spread parameter associated with the distribution over ω , and this may be set using previous knowledge about the expected value of the measurement x_0 , and equation 6 assuming no previous data points have been taken.

To make this concrete, we will determine the prior distribution parameters ν and v used in all segments of the data analysis below. We first set $\nu=1$ as recommended by Mendel. We then note that the error can only take on values in the range $[0-\pi]$. Assuming a uniform distribution over a sphere for errors and calculating the expected value of the squared error yields,

$$E(z^2) = \frac{\int_0^\pi \theta^2 (2\pi \sin(\theta)) d\theta}{\int_0^\pi 2\pi \sin(\theta) d\theta}$$

$$= \frac{\pi^2 - 4}{2} \tag{F.4}$$

To determine v , we assume $d=3$, $\nu'=\nu=1$ (no additional data), set equation (F.4) equal to the left hand side of equation (F.3), and solve for $v=v'$ to get

$$v \sim 0.98$$

Bibliography

1. Abbott, Kathy H. "Exploration of Expert Systems Concepts for Onboard Diagnosis of Faults in a Turbofan Aircraft Engine." Presented at the *1985 American Control Conference*. June 19-21 1985. Boston, Massachusetts.
2. Agu, Masahir, and Yamanaka, Kazuo. "Bayesian estimation of nonlinear Markovian Variables by Partial Observation." *Physical Review A*. Volume 28, Number 6. December, 1983. pp. 3664-3665.
3. Amram, Joseph A. "Decoding and Control Procedures for Partially Observable Markov Systems." M.I.T. Sc.D. Theses, Cambridge, Ma., 1982.
4. Bahill, A. Terry, and McDonald, Jack D. "Model Emulates Human Smooth Pursuit System Producing Zero-Latency Target Tracking." *Biological Cybernetics* 48. pp 213-222.
5. Barnett, Barbara, and Wickens, Christopher D. "Non-optimality in the Diagnosis of Dynamic System States." Technical Report No. CPL 86-8, MDA903-83-K-0255. October, 1986.
6. Baron, Sheldon. "A Control Theoretic Approach to Modelling Human Supervisory Control of Dynamic Systems." *Advances in Man-Machine Systems Research*, Volume 1, pp. 1-47.
7. Boettcher, Kevin L., and Tenney, Robert R. "A Procedure-Based Approach to Human Information Processing." LIDS-P-1499. Laboratory for Information and Decision Systems. Massachusetts Institute of Technology. Cambridge, Massachusetts.
8. Borys, B.B., Johannsen, H.G., Hansel, H.G, and Schmidt, J. "Task and Knowledge Analysis in Coal-fired Power Plants." *Proc. Symp. Human-Computer Interaction and Cognitive Engineering*. IEEE Intern. Conf. Systems, Man, Cybernetics., Atlanta Georgia, October, 1986.
9. Braune, R.J., and Trollip, S.R. "Towards an Internal Model in Pilot Training." *Aviation, Space, and Environmental Medicine*. October, 1982. pp. 996-999.
10. Charny, Leonid "A Decision Support System for Multiattribute Decision Making" *Man-Machine Systems Laboratory Internal Lab Memo*. Massachusetts Institute of Technology. 1987.
11. Cochran, John E. Jr., and Lahr, Brian S. "Satellite Recovery -- Attitude Dynamics of the Targets." In "Advances in the Astronautical Sciences," Kaufman, Bernard, et.al., editors. pp 157-180. Univelt, Inc. San Diego, California. 1986.
12. Conant, Roger C. and Ashby, W. Ross. "Every good regulator of a system must be a model of that system." *International Journal of Systems Sci.*, 1970, Volume 1, Number 2. pp 89-97.

13. de Kleer, Johan; and Brown, John Seely. "Assumptions and Ambiguities in Mechanistic Mental Models." Cognitive and Instructional Sciences Series CIS-9. Xerox Palo Alto Research Center. Palo Alto, California. 1982.
14. du Val, Patrick. "Homographies, Quaternions, and Rotations." Oxford University Press, Amen House, London, E.C.4. 1964.
15. Edwards, Ward, Lindman, Harold, and Savage, Leonard J. "Bayesian Statistical Inference for Psychological Research." *Psychological Review*. Volume 70, Number 3. May 1963. pp 193-242.
16. Erlick, Dwight E. "Judgements of the Relative Frequency of Two Random Sequential Events: Effects of Duration of Observations." WADD Technical Report 60-673. Office of Technical Services, U.S. Department of Commerce, Washington D.C.
17. Fragola, Joseph R. "A Resource Limited Model for Operator Response to Multiple Contemporaneous Events." *Transactions of the American Nuclear Society*. Volume 45, TANSO 45 1-884 (1983).
18. Freedy, Amos; Davis, Kent B. et. al. "*Adaptive Computer Aiding in Dynamic Decision Processes: Methodology, Evaluation, and Applications*." PFTR-1016-76-8/30, August, 1976.
19. Garcia, Carlos E., and Morari, Manfred. "Internal Model Control. 1. A Unifying Review and Some New Results." *Ind. Eng. Chem. Process Des. Dev.* 1982, Volume 21. pp 208-323.
20. Gelb, Arthur, ed. "Applied Optimal Estimation." M.I.T. Press, Cambridge, Massachusetts. 1974.
21. Greenstein, Joel S.; and Rouse, William B. "A Model of Human Decisionmaking in Multiple Process Monitoring Situations." *IEEE Transactions on Systems, Man and Cybernetics*. Volume SMC-12, Number 2, March/April, 1982. pp. 182-193.
22. Hartley, Craig S.; Cwynar, David J.; Garcia, Kathy D.; and Schein, Robert A. "Capture of Satellites Having Rotational Motion." In *Proceedings of the Human Factors Society 30th Annual Meeting*. 1986.
23. Henneman, Richard I. "Human Problem Solving in Complex Hierarchical Large Scale Systems." Center for Man-Machine Systems Research Report Number 85-1. Georgia Institute of Technology. Atlanta, Georgia.
24. Jagacinski, R.J. "Describing the Human Operator's Internal Model of a Dynamic System." *Human Factors*, Vol. 20, No. 4, pp. 425-433, 1978.
25. Jagacinski, R.J., Burke, Micheal W., and Miller, Dwight P. "Use of Schemata and Acceleration Information in Stopping a Pendulumlike System." *Journal of Experimental Psychology: Human Perception and Performance*. 1977, Volume 3, Number 2. pp. 212-223.

26. Johnson-Laird, P.N. "Mental Models in Cognitive Science." *Cognitive Science*. Volume 4, 1980. pp 71-115.
27. Junkins, John L., and Turner, James D. "Optimal Spacecraft Rotational Maneuvers." Elsevier, New York, NY. 1986.
28. Kakalick, James Steven, "Optimum Policies for Partially Observable Markov Systems." NR 276-004 #18, 1976.
29. Kieras, David E., and Bovair, Susan. "The Role of a Mental Model in Learning to Operate a Device." *Cognitive Science* 8, 255-273 (1984).
30. Kleinman, David L., Pattipati, Krishna R., and Ephrath, Arye R. "Quantifying an Internal Model of Target Motion in a Manual Tracking Task." *IEEE Transactions on Systems, Man, and Cybernetics*. Volume SMC-10, Number 10, October, 1986. pp. 624-636.
31. Kleinman, David L., and Curry, Renwick E. "Some New Control Theoretic Models for Human Operator Display Monitoring." *IEEE Transactions on Systems, Man, and Cybernetics*. Volume SMC-7, Number 11, November, 1977. pp. 650-651.
32. Kleinman, David L., and Curry, Renwick E. "An Equivalence Between Two Representations for Human Attention Sharing." *IEEE Transactions on Systems, Man, and Cybernetics*. September, 1976. pp. 650-651.
33. Laios, L. "Predictive Aids for Discrete Decision Tasks with Input Uncertainty," *IEEE Transactions on Systems, Man and Cybernetics*. Vol. SMC-8, No. 1, January, 1978. pp 19-29.
34. Laritz, Frank, and Sheridan, T.B. "Evaluation of Fuzzy Rulemaking for Expert Systems for Failure Detection." Man-Machine Systems Lab Report. 1984.
35. Larkin, Jill. "Cognition of Learning Physics." *Am. J. Phys.* Volume 49, Number 6, June 1981. pp 534-541.
36. Lehner, Paul E.; Rook, Frederick W.; and Adelman, Leonard. "Mental Models and Cooperative Problem Solving with Expert Systems." PAR Report No. 84-116. September 1984.
37. Marra, Antonio Jr. "Human Operator Performance in Dynamic Multi-Task Environments." Master's Thesis in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology, Cambridge, MA. 1988.
38. Mendel, Max "Optimal Combination of Information from Various Sources." Internal Lab Memo, Man-Machine Systems Lab. Massachusetts Institute of Technology, Cambridge, MA. 1987.

39. Mendel, Max "Prediction of Discrete Time Series: A Bayesian Approach." Internal Lab Memo, Man-Machine Systems Lab. Massachusetts Institute of Technology, Cambridge, MA. 1987.
40. Miller, Duncan C.; and Elkind, Jerome I. "The Adaptive Response of the Human Controller to Sudden Changes in Controlled Process Dynamics." *IEEE Transactions on Human Factors*. Vol. HFE-8, No. 3, September 1967. pp 218-223.
41. Niemela, R.J.; and Krendel, E.S. "Detection of a Change in Plant Dynamics in a Man-Machine System," *IEEE Transactions on Systems, Man and Cybernetics*. November, 1975.
42. Palmer, James D. "Large Scale Systems: Systems, Man, and Cybernetics Overview." *IEEE Transactions on Systems, Man, and Cybernetics*. Volume SMC-13, Number 4. July/August 1983. pp. 453-459.
43. Pinker, Steven. "Emergent Two-Dimensional Patterns in Images Rotated in Depth." *Journal of Experimental Psychology: Human Perception and Performance*. 1980, Volume 6, Number 2. pp. 244-264.
44. Rasmussen, J. and Goodstein, L.P. "Decision Support in Supervisory Control of High-risk Industrial Systems." Proceedings of the 2nd IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-Machine Systems. September, 1985. Proceedings published by Pergamon Books Limited, U.K.
45. Rasmussen, J. "Skills, rules, and knowledge; signals, signs, and symbols; and other distinctions in human performance." *IEEE Transactions on Systems, Man, and Cybernetics*. SMC-13, 257-266. 1983.
46. Rasmussen, J. "Outlines of a Hybrid Model of the Process Plant Operator." In T. Sheridan, & G. Johannsen (Eds.), *Monitoring Behavior and Supervisory Control*, New York, N.Y.: Plenum Press, 1976.
47. Rice, John R.; Yorchak, John P.; and Hartley, Craig S. "Planning for Unanticipated Satellite Servicing Operations." In *Proceedings of the Human Factors Society 30th Annual Meeting*. 1986.
48. Rieger, Chuck; and Grinberg, Milt. "The Causal Representation and Simulation of Physical Mechanisms." Technical Report TR-495. N00014-76C-0477. November, 1976.
49. Rouse, William B., and Rouse, Sandra H. "Analysis and Classification of Human Error." *IEEE Transactions on Systems, Man, and Cybernetics*. Volume SMC-13. Number 4. July/August 1983. pp 539-549.
50. Rouse, William B. "A Model of Human Decisionmaking in Stochastic Estimation Tasks," *IEEE Transactions on Systems, Man and Cybernetics*. Vol. SMC-7, No. 4, April, 1977.

51. Sage, Andrew P. "Behavioral and Organizational Considerations in the Design of Information Systems and Processes for Planning and Decision Support," *IEEE Transactions on Systems, Man, and Cybernetics*. Vol. SMC-11, No. 9, September, 1981.
52. Sheridan, T.B. "The System Perspective." In *Human Factors in Modern Aviation*, E.L. Weiner and D.C. Nagel, Eds. 1985.
53. Sheridan, T.B. "Interaction of Human Cognitive Models and Computer-based Models in Supervisory Control." Man-Machine Systems Laboratory Report. Massachusetts Institute of Technology, Cambridge, Massachusetts. March 1984.
54. Sheridan, T.B.; and Ferrell, William R.; "Man-Machine Systems." The M.I.T. Press. Cambridge, Massachusetts. 1974.
55. Sheridan, T.B., Chairman; Jenkins, James P.; and Kisner, Roger P. "Proceedings of the Workshop on Cognitive Modelling of Nuclear Plant Control Room Operators." NUREG/CR-3114. ORNL/TM-8614. Dedham, Massachusetts. August 15-18, 1982.
56. Sheridan, T.B. "Optimal Allocation of Personal Presence." *IEEE Transactions on Systems, Man and Cybernetics*. Volume SSC-6, Number 3. July, 1970. pp 242-244.
57. Smallwood, R.D. "Internal Models and the Human Instrument Monitor," *IEEE Trans. on Human Factors in Elec.* Vol. 8, No. 3, pp.181-187, September 1967.
58. Stevens, A.L.; and Collins, Allan. Multiple Conceptual Models of a Complex System." In R.N. Snow, P. Federico, & W.E. Montague (Eds.), *Aptitude, Learning, and Instruction*. Hillsdale, N.J., Lawrence Erlbaum Associates, 1980.
59. Tanabe, T., Koyama, H., and Ohyama, E. "Model Based Vision System for Autonomous Teleoperated Spacecraft." Joint AAS/Japanese Rocket Society (JRS) Symposium. December 15-19, 1985.
60. Thijs, Wim L.Th.; Stassen, Hank G.; Kok, Jan J.; and van der Veldt, Rolf J. "Supervisory Control and Fault Management. Voyage of Discovery Through the Estate of Theory and Model Creation." Preprint for presentation at the *IX Triennial IFAC World Congress*. Budapest, Hungary, July 2-6, 1984.
61. Tsach, Uri.; Sheridan, Thomas B.; and Buharali, Ahmet. "Failure Detection and Location in Process Control - Integrating a New Model Based Technique with Other Methods." Presented at the *1983 American Control Conference*. June 22-24, San Francisco, CA.
62. Tsach, Uri. "Failure Detection and Location Method (FDLM)." Man-Machine Systems Laboratory Technical Report. Massachusetts Institute of Technology. Cambridge, MA. Sept. 30, 1982.

63. Tulga, M. Kamil. "Dynamic Decision Making in Multi-Task Supervisory Control: Comparison of an Optimal Algorithm to Human Behavior." Sc.D. Thesis in the Department of Mechanical Engineering at the Massachusetts Institute of Technology, Cambridge, Massachusetts. 1978.
64. van Bussel, F.J.J. "Human Prediction of Time Series," *IEEE Transactions on Systems, Man and Cybernetics*. Vol. SMC-10, No. 7, July, 1980.
65. van Heusden, Arnold R. "Human Prediction of Third-Order Autoregressive Time Series," *IEEE Transactions on Systems, Man and Cybernetics*. Vol. SMC-10, No. 1, January, 1980.
66. Veldhuyzen, and Stassen, Hank G. "The Internal Model Concept: An Application to Modeling Human Control of Large Ships." *Human Factors*. Volume 19, Number 4. 1977. pp. 367-380.
67. Waltar, Alan E, session organizer. "Safety Implications of Man/Machine Interface Developments." *American Nuclear Society Transactions*. Volume 45, 1983. pp.357-363.
68. Wertz, James R., ed. "Spacedraft Attitude Determination and Control." D. Reidel Publishing Company, Dordrecht, Holland. 1985.
69. Wickens, Christopher D. "An Investigation of the Properties of the Operator's Internal Model of Complex Dynamic Systems." Technical Report No. CPL 86-9, MDA903-83-K-0255. October 1986.
70. Young, Richard M. "The Machine Inside the Machine: Users' Models of Pocket Calculators." *Int. J. Man-Machine Studies*. Volume 15, 1985. pp 51-85.