N79-17498

INTERRUPTED MONITORING OF A STOCHASTIC PROCESS

by Everett Palmer

Ames Research Center, NASA, Moffett Field, CA. 94035

ABSTRACT

As computers are added to the cockpit, the pilot's job is changing from one of manually flying the aircraft, to one of supervising computers which are doing navigation, guidance and energy management calculations as well as automatically flying the aircraft. In this supervisorial role the pilot must divide his attention between monitoring the aircraft's performance and giving commands to the computer. In this paper, normative strategies are developed for tasks where the pilot must interrupt his monitoring of a stochastic process in order to attend to other duties. Results are given as to how characteristics of the stochastic process and the other tasks affect the optimal strategies. The optimum strategy is also compared to the strategies used by subjects in a pilot experiment.

INTRODUCTION

"New York control, this is NASA 1 arriving on CARMEL 1 with an expected arrival time at MERGE waypoint of 14:31:04." "NASA 1, you are cleared to arrive on CARMEL 1, with a merge time of 14:32:18." This exchange between pilot and controller occured in a recent Ames simulation study of 4D RNAV in the terminal area[1,2]. The pilot was cleared for a different RNAV approach route and arrival time. The pilot next entered this data into his onboard navigation and guidance computer. In doing this he had to divide his attention between monitoring the autopilot's performance with his flight instruments and entering data into the computer through his multifunction display and keyboard. Observations of how pilots divided their attention between monitoring and data entry tasks were the motivation for the modeling and more structured experimental work on attention sharing presented in this paper.

The environment in which the pilot interacts with his onboard computer is quite different from other jobs where a person interacts with a computer. In a management information system, teleoperator control, or in most human interaction with a computer, the computer is, or can easily be halted to allow the person time to think and plan his next input. The person and the computer work sequentially. When an aircraft is being controlled in real time by a computer it can not be stopped while the pilot

leisurely inputs his commands. In this environment both computer and man must work in parallel. The pilot must interrupt his monitoring to interact with the computer. He must also interrupt the discrete tasks to monitor. Other characteristics of discrete tasks and monitoring in the cockpit are the following. The discrete tasks are presented at random. They should be accomplished by a certain time but usually sufficient time is available to do the tasks. Attention must be diverted from monitoring for fairly long blocks of time (seconds) to do the discrete tasks. The displays the pilot must monitor show the error between his vehicle's state and the desired state. When the aircraft is controlled by an autopilot, these signals are relatively low bandwidth signals that should be monitored for out of tolerance readings.

The objective of this research is to determine how design parameters of both displays and the computer interface affect monitoring and data entry performance. In this paper, a task is developed which has many of the above characteristics but which is simple enough so that the attention allocation problem has an optimal solution. A model based on the internal model concept is developed for this task. The model treats the discrete tasks as a constraint and then uses a dynamic programming formulation to maximize monitoring performance subject to the constraint of finishing the discrete tasks on time. Results are presented as to how well pilots could monitor a first order process for out of tolerance signals when they were allowed to monitor the display only half the time.

SPECIFIC PROBLEM

Process Dynamics: The process to be monitored is the output of a first order filter driven by white gaussian noise. The display (fig.1) is updated every 2 seconds and is quantized into 11 cells — .50 σ wide. The display is defined as being out of tolerance if it is in the outermost 2 cells ($|x| > 1.75$ σ). In figure 1 these states are indicated with + signs. The process bandwidth determines how predictable the signal is. The ratio of the tolerance to the output variance determines how frequently the signal will be out of tolerance.

Monitoring Task: Whenever the subject observes the process as being out of tolerance he gets a reward of one unit.

Discrete Task: At each time at which the display is observed, the subject decides to either monitor next time or to divert his attention to the discrete task for one or more units of time. The subject was constrained to do m discrete tasks in the next n time units. The objective was to maximize the reward for monitoring subject to the constraint of finishing the discrete tasks. This constraint formulation seems to be a good description of the real situation. It does not require the experimenter to specify a cost function stating the relative worth of time spent on monitoring and discrete tasks.

237

THEORY

A review of the literature in the fields of manual control, human factors and psychology found a number of empirical studies which required the operator to interrupt monitoring tasks to do discrete tasks. Models have also been developed for either instrument monitoring or discrete tasks. No papers were found which addressed the problem of what strategies operators use, or should use, to time share their attention between monitoring and discrete tasks. However, Smallwood's paper (3) on human instrument monitoring proposes an approach which can be applied to the present problem. This approach makes the reasonable assumption that the operator has an internal model of the process he is monitoring and of the environmental factors that affect the process. This internal model can be used to predict the future behavior of the process. Smallwood makes the following assumptions that describe how the operator reacts to environmental inputs.

Assumption 1: The human operator bases his state of information about his environment upon an internal model of this environment; the model is formed as a result of past perceptions of his environment.

Assumption 2: The human operator behaves optimally with respect to his task and his correct state of information within his psycho-physical limitations.

The structure of this model is shown in figure 2. The key problems in using this approach are to discover the form of the operator's internal model and the optimal response. If the operator's model of the process is exact and he has no psycho-physical limitations the resulting model is normative. Introducing errors in the internal model and psycho-physical limitations such as observation noise, reaction time or faulty memory converts the original normative model into a descriptive model of human behavior.

In the following models, it is assumed that the operator's internal model of the process and environmental disturbances is exact. He knows the parameters of the process and can use this knowledge to predict the probability of being in a particular state in k seconds given he knows the current state. For a first order process with bandwidth w, the distribution of the position of the display after last observing the display t seconds ago at position x_0 is a gaussian distribution with

mean $m(t) = x_0 e^{-w t}$

variance $v(t) = \sigma_0^2 (1-e^{-2 w t})$

Figure 3 plots these distributions and the probability that the signal will be out of tolerance in the future for various values of x_0.

A decision is made after each monitoring observation of how many stages to devote to discrete tasks. The decision may be to do no discrete tasks in which case the operator continues to monitor.
Define

$P_{ij}(k)$= the probability that the process will be in state j in k stages, given the process if in state i.

$f_n(m,i)$= the maximum expected return when the process is observed in state i with n stages to go and m discrete tasks remain to be done.

q= the number of stages devoted to discrete tasks before the next monitoring observation when the process is observed in state i with n stages to go.

C= set up cost, the number of stages wasted when attention is shifted to discrete tasks.

$q^* = q-C$ if $q-C \geq 0$
$q = 0$ if $q-C < 0$

then

$$f_n(m,i) = \max_{0 \leq q \leq m} \left[\sum_{jout} P_{ij}(q) + \sum_{all\ j} P_{ij}(q) f_{n-q-1}(m-q^*,j) \right]$$

$$f_n(m,i) = 0 \text{ if } n=m$$

Figure 4 is a cross plot of two measures of the optimal strategy 4 for different values of process bandwidth. The measures are the fraction of time spent doing discrete tasks or not monitoring the display, f(tasks)=m/n, and the fraction of observed out of tolerance signals to the total number of out of tolerance signals, f(hits)= $\Sigma f_n(m,i)/n$. The figure shows that if the optimal strategy is followed when monitoring a first order process with a bandwidth of .2 rad/stage, 85% of the out of tolerance signals will be observed even if only 50% of the time is spent monitoring. Figure 4 also shows the monitoring performance that would be expected if the pilot could make perfect predictions and his expected performance if he could make no predictions. As the bandwidth of the process decreases and the signal becomes more predictable performance approaches that possible with perfect information.

In many discrete tasks there is the equivalent of a set up cost each time the task is started or restarted after being interrupted. For example, in entering data into a keyboard, some time is lost while the pilot shifts his attention to the keyboard and positions his hands. Figures 5 and 6 show the effect of a discrete task set up cost on sampling strategy and monitoring performance. As the set up cost increases, the best strategy is to look away for longer and longer periods of time when the display is observed near the center. With a set up cost of 2, if the display is observed in the center, the best policy is to complete all of the discrete tasks with no interruption. This is

why the monitoring performance shown in figure 6 for a set up cost of 2 is so close to the performance that is possible when no predictions are made.

Figure 7 shows the sensitivity of monitoring performance to discrete task chunk size - the minimum number of stages which must be spent on discrete tasks. Note that when the minimum chunk size is 5 that the decrement in performance is only large when less than 50% of the time must be spent on discrete tasks. This is because when f(tasks) is greater than 50% the optimum strategy is to look away for more than 5 stages so that a minimum chunk size of 5 is only a minor constraint on performance.

Finally figure 8 shows the sensitivity of monitoring performance to display tolerance or the probability that the display will be out of tolerance.

EXPERIMENT

Results are presented in this section as to how well subjects could monitor a first order process for out of tolerance signals when they were allowed to devote only half of their time to monitoring. The objective of this exploratory study was to determine what strategy subjects used and how it differed from the optimal strategy. To this end, the monitoring task was designed to be identical to the monitoring problem solved by the dynamic programming model. It was not expected that the subjects would behave optimally but a number of features of the optimal strategy seem to be fairly intuitive. Namely it is best to look away longer when the display is farthest from the out of tolerance boundary, look away longer and less frequently with a setup cost, and look away longer if the ratio of discrete tasks to go approaches 1.

The subjects monitored the output of a first order filter (fig.2 rad/stage) driven by white noise. The display, shown in figure 1, was quantized into 11 states, .50 apart. The display states were numbered from left to right with state 6 in the center. Since the display is completely symetric about state 6, the strategy in states 5&7,4&9,.....1&11 are identical. In the following discussion we will refer to only the left side of the display. The process was defined as being out of tolerance whenever it was in the outermost two states. The signal was out of tolerance 8.0 percent of the time.

The CRT terminal (figure 1) showed two integers to the left of the monitoring display. The left integer showed the number of stages remaining in which the subject had to finish the discrete tasks. The right integer showed the number of discrete tasks left to do (m). The display was updated every 2 seconds. On each update the number of stages to go (n) was decremented by 1. In order to do a discrete task, the subject pushed the space bar on the keyboard and at the next display update only the two integers were displayed and both were decremented by 1. To switch

back to the monitoring display the subject pushed the space bar again.

The subjects were instructed to attempt to conserve as many out of tolerance signals as possible subject to the constraint of finishing all of the discrete tasks (m) before the number of stages to go (n) reached zero. On some of the runs a set up cost of 1 stage was added to the discrete task. In this case if the subject switched to discrete tasks for k stages, the number of stages to go was decremented by k and the number of tasks to do was decremented by k-1.

One replication of this experiment consisted of 42 stages. The number of stages to go (n) and the number of discrete tasks to do (m) were initially set to 40 and 24 respectively. When n reached 0, it was reset to 40 and m was reset to 24 and the next replication began.

Four airline pilots served as subjects in this experiment. Each subject monitored 7 blocks of 6 replications each for a total of 1680 stages (7x6x40). During the first block the subject just monitored the process dynamics. On the next three blocks there were no discrete tasks to do on the first replication but the subject had to look away from the display for 20 stages on each of the remaining 5 replications. On the last 3 blocks a set up cost of 1 stage was introduced.

Subjects were given a 5 minute rest between each block. Each subject monitored the same random process as the other subjects.

For each replication, performance measures included the number of times the display was out and observed out (a hit), the subjects sampling strategy at each stage, and the optimal strategy at each stage. The above data was also collected for an "autopilot" following the optimal strategy for the same random sequence monitored by the subjects.

At the end of the experiment the subjects completed a questionaire in which they were asked questions about the strategy they used and to rate the experimental tasks on a set of semantic differential scales.

RESULTS

Figure 9 shows the fraction of hits for the four subjects, the autopilot,and the model. The "autopilot" results are the optimum performance for the particular random sequence used in the experiment. The model results are the expected performance for an infinitely long random sequence. As expected no subject did as well as the autopilot, however the subjects monitoring performance with no set up cost was well above the chance value of .5. The set up cost however caused a large decrement in per-

259

on the other hand tended to finish early. This type of behavior are if at the beginning of each replication they had looked away for 21 stages and completed the discrete tasks all at once.

With a set up cost the subjects should have looked away less often and for more stages than they actually did. Figure 1a shows the frequency of decisions of various lengths for the subjects and the autopilot. With no setup cost the subjects average decision was 1.8 stages whereas the autopilot's average decision was only slightly longer at 2.1 stages. The main difference between the subject and autopilot was that the subjects looked away for only 1 stage almost twice as often as the autopilot.

With the set up cost both subjects and autopilot increased the average lengths of their decisions to 3.6 and 4.2 respectively. Note however that the subjects diverted attention to discrete tasks almost twice as often as the autopilot (7.6 vs 4.4). The subject's also made a number of decisions of length 1 which with a set up cost of 1 accomplished nothing.

The optimal strategy is always to look away from the monitoring display for more stages when the display is observed near its center (state 6). The data in Table 1 is similar to that in figure 10 but table 1 also shows the effect of display state on subject and autopilot sampling decisions. The subject's decisions were a strong function of the display state. The subjects almost never looked away from the display when it was out of tolerance and only a few times when it was almost out (state 3). The closer the display was to its center, the longer the subjects looked away - in general agreement with the optimal strategy. Note however that in addition to not looking away quite as long as the autopilot, that on a number of occasions the subjects continued to monitor when the display was near the center (states 5 or 6). The autopilot always looked away in these states if there were any discrete tasks remaining to be done. Some of the subjects decisions to continue monitoring were probably due to the forced pace -real time nature of the task because if the subject failed to push the space bar he would continue to monitor by default. This conjecture is partially reinforced by the pilots description of their strategy. All pilots reported that their strategy was to look away for various numbers of stages when the display was near its center.

Figure 11 plots the number of discrete tasks left to do (m) vs. the number of stages to go (n) for the subjects and the autopilot. It shows whether the discrete tasks were done early or late in the replication. With no set up cost the subjects and the autopilot behaved similarly- both keeping the ratio of m to n very slightly less than 0.50 and thereby spreading the discrete tasks through out the time available. Note however that both autopilot and subjects tended to finish a few stages before the end of the replication. The right hand graph shows that with a set up cost of 1 the autopilot and subjects behavior was quite different. The subjects did not look away long enough and therefore got behind in finishing the discrete tasks. The autopilot

on the other hand tended to finish early. This type of behavior prevents the autopilot from becoming trapped if the display starts to go out of tolerance with a few stages left to go. The right hand graph indicates that this rational behavior is not what the subjects intuitively did.

Table 2 summarizes the results of comparing each decision made by the subjects with the optimal decision. In states 1,2 and 3 the subjects tend to perform optimally. In state 4, the subjects look away for more stages than is optimal - especially when there is a set up cost. In states 5 and 5 , the subjects did not look away long enough. These observations hold with and with out a setup cost. However subjects made considerably more optimal and near optimal decisions with no set up cost.

During the debriefing, 2 subjects said that with zero set up cost their strategy was to look away for 3 times in state 6, 2 times in state 5, and 1 time in state 4, and n times otherwise. The other two subjects said that they looked away 2 or 3 times in state 5 or 6 and did not look away otherwise. On the average , these strategies will result in monitoring approximately 50 % of the time. With the set up cost, the subjects stated that they looked away longer (typically 4 or 5 stages) although one subject said he did not feel he could observe the display often enough to do anything but guess. Subjects stated that the display state was the key thing that influenced their strategy. The number of stages to go and the number of discrete tasks to do had little effect on their strategy unless they were running out of time and then they stated they would look away for a longer number of stages.

Figure 12 shows the average subject ratings on a set of semantic differential scales. Those adjectives have been ordered so that the ratings for the task with no set up cost are to the right of the ratings with a set up cost. Considering only scales with a difference of one or or more, the task with a set up cost was demanding, hard to learn, confusing ,surprising, annoying, active, complex, and frustrating.

CONCLUDING REMARKS

In this paper the general problem of time sharing attention between monitoring and other duties has been described and a dynamic programing model for attention sharing was presented. Model performance was presented in terms of the fraction of out of tolerance signals seen as a function of the amount of time spent on non-monitoring duties. The effect of such parameters as process bandwidth and tolerance and discrete task set up cost and chunk size on monitoring performance and the normative time sharing strategies was shown. Future work will extend this model to multiple second order processes and incorporate human limitations such as observation noise and risk aversion.

In the experiment subjects monitored the output of a first

order process for out of tolerance signals. The subjects did not perform as well as an autopilot following the optimal strategy for this task. however their time sharing strategies were a strong function of the display state and a weaker function of the ratio of the number of discrete tasks to do divided by the number of stages to go. With a set up cost, the subjects looked away for fewer but longer amounts of time but the optimal strategy required even longer diversions of attention to discrete tasks. The fact that the subjects did not look away as long as was optimal may be attributable to risk averse behavior.

Future experiments will use a continuous version of the monitoring task used in this experiment with second order dynamics. The effect of the discrete task parameters - set up cost, chunk size, time required, and time available - on monitoring performance will be determined. Process variables will include bandwidth and the number of displays. The results of this experiment and the modeling work will be used to predict monitoring performance in an experiment in which subjects are required to divide their attention between an actual data entry task and display monitoring.

REFERENCES

1. L.Tobias,"Simulator Evaluation of Terminal Area RNAV Concepts",NASA TN , In Press.

2. E. Palmer,"Pilot's Manual for the 4D Area Navigation and Auto-pilot Systems in the Flight Management Research Simulator",NASA TMX-73180,October 1975.

3. R.D. Smallwood,"Internal Models and the Human Instrument Monitor", IEEE Trans. on Human Factors in Electronics ,Vol. HFE-8,No. 3, Sept. 1967,pp.181-187.

4. R.A.Howard,"Dynamic Programing and Markov Pr es", MIT Press,1968.

Table 1. Subject and autopilot strategy as a function of display state for set up costs of 0 and 1. The data are the average number of decisions for each display state for 600 stages of monitoring when n>? and n-m>?.
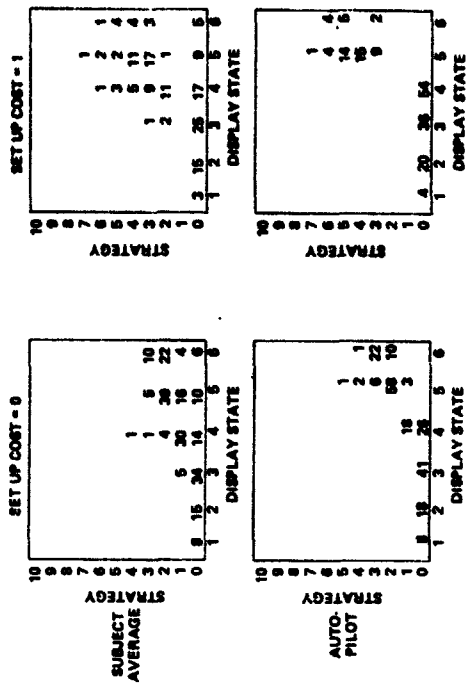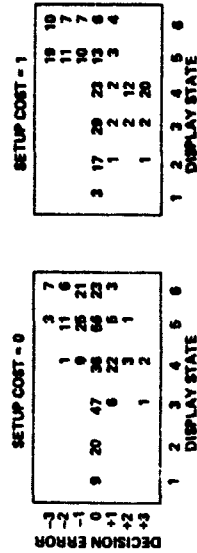

Table 2. Error between the optimum and the subjects decisions as a function of the display state for setup costs of 0 and 1. The data are the average number of decision errors for each display state for 600 stages of monitoring. Decision error = Optimum decision - subject decision.
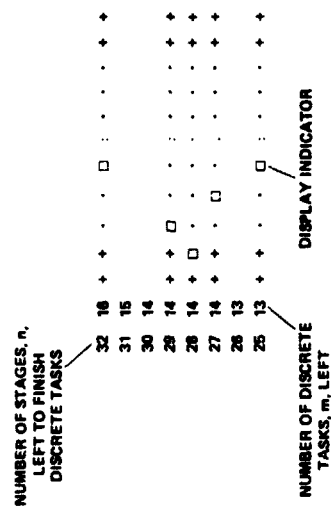
Figure 1. The monitoring display used in the experiment. A new display line was added every 2 seconds. The display was quantized into 11 cells - .59 σ wide. The display was out of tolerance if it was in the outermost 2 cells indicated with the + signs. At stages 22 and 27 this subject decided to look away from the display to do discrete tasks for 2 and 1 stage respectively.
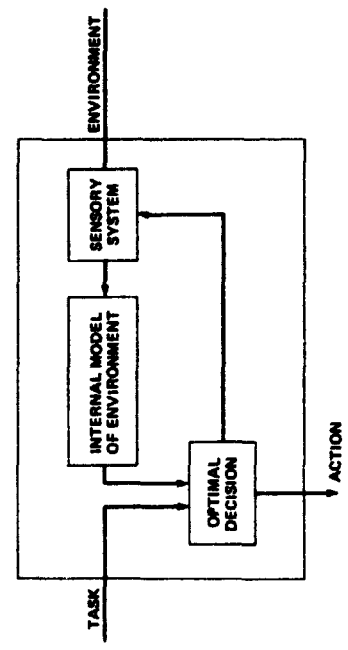
NUMBER OF STAGES, n, LEFT TO FINISH DISCRETE TASKS

NUMBER OF DISCRETE TASKS, m, LEFT

DISPLAY INDICATOR

Figure 2. A block diagram of the human monitor (from Smallwood (3)).

ENVIRONMENT

SENSORY SYSTEM

INTERNAL MODEL OF ENVIRONMENT

OPTIMAL DECISION
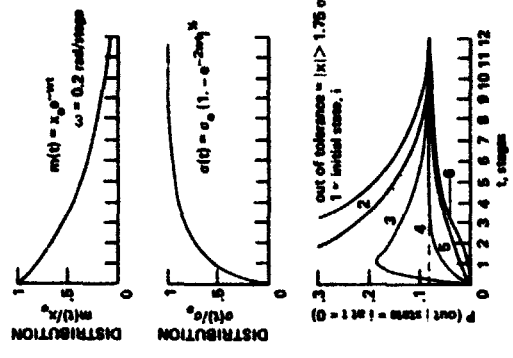
TASK

ACTION

Figure 3. The state of information of a perfect monitor after looking away from the output of a first order filter with bandwidth 0.2 rad/stage driven by white noise.

$m(t) = x_o e^{-\alpha t}$

$\alpha = 0.2$ rad/stage

$\sigma(t) = \sigma_o (1 - e^{-2\alpha t})^{\frac{1}{2}}$

out of tolerance = $|x| > 1.75 \sigma$

i = initial state, i

P (out | state = i at $t = 0$)

DISTRIBUTION $m(t)/x_o$

DISTRIBUTION $\sigma(t)/\sigma_o$

t, stages
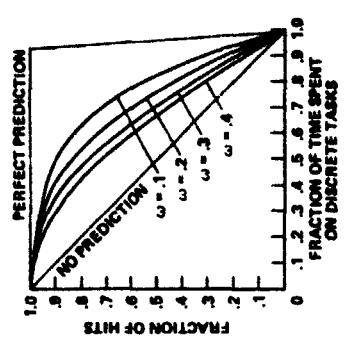
Figure 4. Fraction of hits vs fraction of time on discrete tasks for four values of process bandwidth. ($\sigma = 1.8$, $T = 1.75$, $P(out) = 2.80$)

PERFECT PREDICTION

NO PREDICTION

FRACTION OF HITS

FRACTION OF TIME SPENT ON DISCRETE TASKS

Figure 7. The effect of discrete task chunk size on monitoring performance. (ω =4.2 rad/stage, n=40 stages, σ =1.6,T=1.75, P(out)=0.689)

Figure 8. The effect of display tolerance on monitoring performance. (ω =0.2 rad/stage, σ =1.0, n=40 stages)

TOLERANCE   P(out)
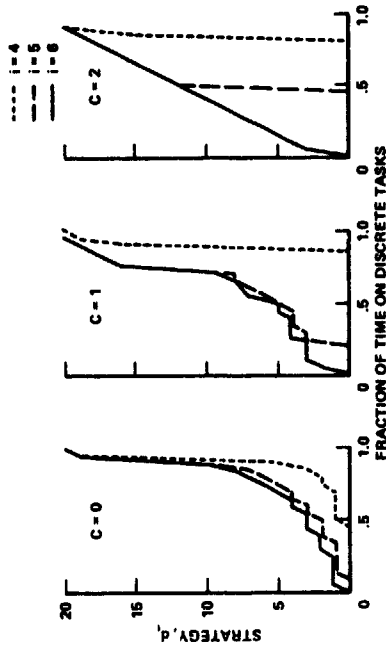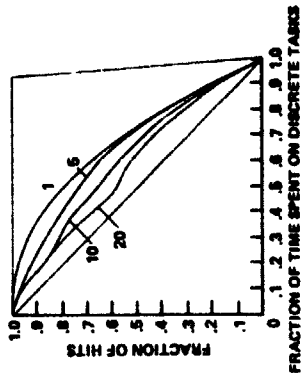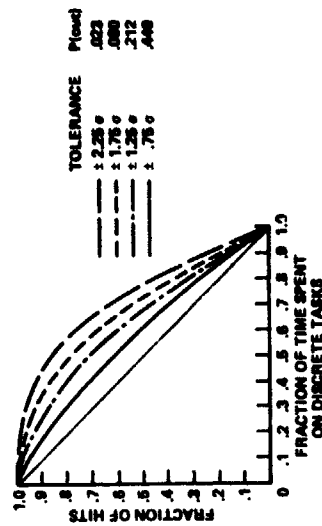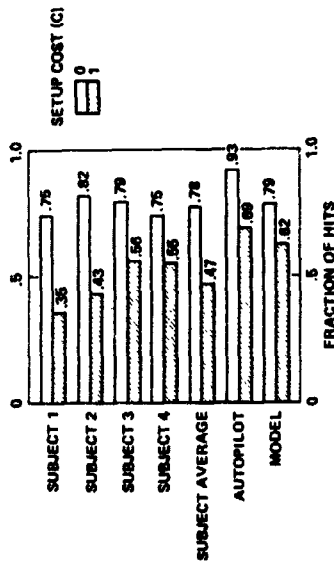± 2.25σ   .023
± 1.75σ   .080
± 1.25σ   .212
± .75σ   .440

Figure 5. The effect of a discrete task set up cost (C) on the optimal time sharing strategy for states 4, 5 and 6. In states 1, 2 and 3 the optimal decision is 0 until the fraction of time which must be devoted to discrete tasks is very high. (n=20 stages,ω=0.2 rad/stage, T=1.75, P(out)=0.994)

Figure 6. The effect of a discrete task set up cost on monitoring performance for the dynamic programing model with discrete task constraint. (n=30 stages, ω =0.2 rad/stage, T=1.75, σ =1.0, P(out)=0.???)

Figure 9. The fraction of hits for discrete task set up costs of 0 and 1.

Figure 10. Frequency of decisions to spend various numbers of stages on discrete tasks for the subjects and the autopilot for set up costs of 0 and 1.
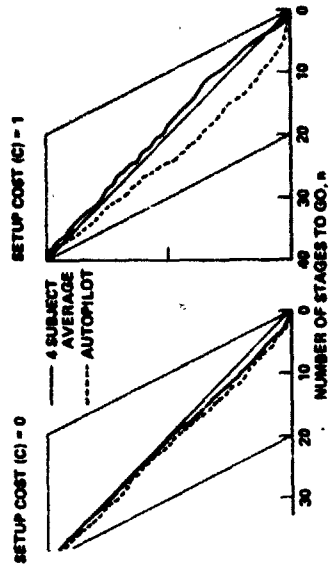

Figure 11. The number of discrete tasks to do (m) vs the number of stages to go (n) for discrete task set up costs of 0 and 1.
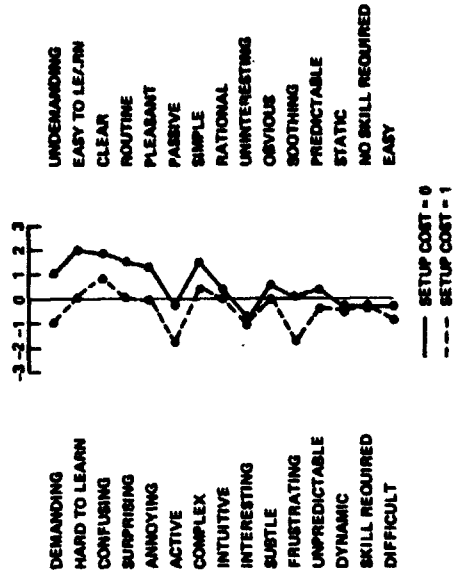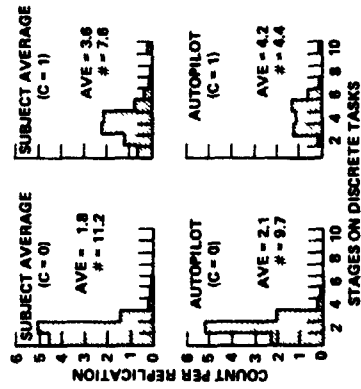

Figure 12. Semantic differential ratings for the experimental task with set up cost of 0 and 1.