EXPERIMENTAL EVALUATION OF THE CONCEPT OF SUPERVISORY MANIPULATION

T.L. Brooks
Robotics and Teleoperator Group
Jet Propulsion Laboratory
California Institute of Technology

T.B. Sheridan
Man-Machine System Laboratory
Massachusetts Institute of Technology

Abstract

A computer-controlled teleoperator system which is based on task-referenced sensor-aided control has been developed to study supervisory manipulation. This system, called SUPERMAN, is capable of performing complicated tasks in real-time by utilizing the operator for high-level functions related to the unpredictable portions of a task, while the subordinate machine performs the more well-defined subtasks under human supervison.

To determine whether supervisory control schemes such as these offer any advantage over manual control under real-time conditions, a number of experiments involving both simple and complicated tasks were performed. Six representative tasks were chosen for the study: (1) obtaining a tool from a rack, (2) returning the tool to the rack, (3) removing a nut, (4) placing samples in a storage bin, (5) opening and closing a valve, and (6) digging with a shovel. The experiments were performed under simulated conditions using four forms of manual control (i.e., switch rate, joystick rate, master-slave position control, and master-slave with force feedback), as well as supervisory control. Through these experiments the effectiveness and quality of control were evaluated on the basis of the time required to complete each portion of the task and the type and number of errors which occurred.

Even under the "best" control conditions (i.e., no degraded sensor or control loops due to time delays, restricted bandwidths, etc.) supervisory control was found to improve performance for all forms of manual control except force-reflecting master-slave which was found to be slightly faster than supervisory control, b more prone to errors. With degraded sensor or control loops it is fairly predictable that supervisory control will show even more advantage, through the latter experiments are yet to be done.

1. INTRODUCTION

Teleoperators have traditionally relied on relatively simple and direct man-machine interfaces for control. However, with the advent of microcomputers and advanced sensor techniques it is now possible to design and build a hierarchical control system in which the operator is responsible for the higher-level functions related to the unpredictable portions of the task, while the subordinate machine performs the more well-defined subtasks under human supervision. Control based on a supervisor-subordinate relationship such as this is called "supervisory control" [1]. In general, the human operator communicates intermittently with the computer, and the computer, in turn and continuously in time, controls the sensors and actuators of the vehicle and manipulator. In essence, the teleoperator system acts as an autonomous "robot" for short periods while in the pursuit of task goals previously programmed by the operator or updated on the last cycle of communication. This mode of control promises more precision for certain tasks, less susceptibility to failure in the event of communication channel breakdown, and greater efficiency than direct human control.

To investigate the relative merits of supervisory control applied to teleoperators, a task-referenced sensor-aided supervisory system, called SUPERMAN, was built and experiments were performed. This paper will evaluate those experiments and by comparison of various conventional control modes with supervisory control, demonstrate that supervisory manipulation does improve performance in the majority of cases.

2. METHOD

Apparatus

The major elements of the SUPERMAN system are a modified Argonne E2 master-slave manipulator with six degrees-of-freedom, a dedicated control interface (DASI), and an Interdata 70 computer. Designed for efficient man-machine interaction with both analog and symbolic control inputs, the system can be commanded by a variety of conventional control modes as well as supervisory. In addition, time delay and/or noise can be added for experimental purposes.

Using both analog and symbolic commands, a manipulation can be taught and/or demonstrated to the computer. Trained manipulations can be transformed from one coordinate system to another so that once the generic characteristics of a task have been learned, the machine can perform similar tasks in different locations without further training. When the human operator requires a particular trained manipulation he simply "initializes" the new coordinate system relative to the old by moving the teleoperator hand to the starting point of the task (e.g., grasping a nut or valve handle) and signals for execution. Certain objects in the task environment can, of course, maintain their original coordinates. For a complete development of task transformations related to supervisory control see refs. 2 and 3.

Since the manipulator can sense the forces generated during the task, supervisory programs can call for repeated movements which, upon certain touch conditions becoming true, branch into other movements. For example, repeated hand movements can grasp a nut, unscrew it by one revolution, pull back to test whether it is off and, if it is, place it in a bucket or, if it is not, repeat the operation. Similar supervisory programs have been applied to attaching a nut to a bolt, opening and closing a valve, scooping dirt and so on. Further information on the SUPERMAN system can be found in ref. 3.

The manipulator laboratory was arranged as shown in fig. 1 during the experiments. To simulate remote conditions the operator viewed the task environment through either a mono or 2-view television system. The video system consisted of two black and white high-resolution 9 in. monitors, a fixed camera with wide angle lens, and a zoom camera with pan & tilt.

Figure 2 shows the manipulator environment and the experimental tasks designed for this study. The tool rack and sample buckets remained in the locations shown throughout the experiments since these pieces of equipment are usually rigidly attached to the teleoperator vehicle in real applications. Also shown in the figure are the movable task hub and task board on which representative tasks such as valves, bolts, etc were mounted. The location of the task hub and board were changed throughout the study to simulate the random task/vehicle relationships which are typical of the arbitrary environments found in marine and space applications.
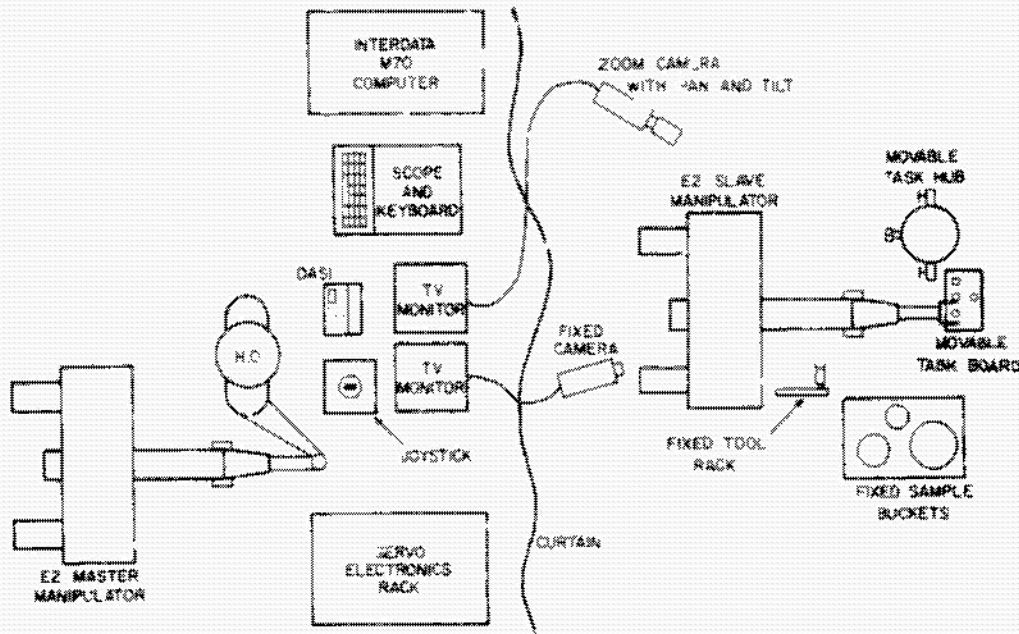
594

Fig. 1. Schematic of Experimental Layout
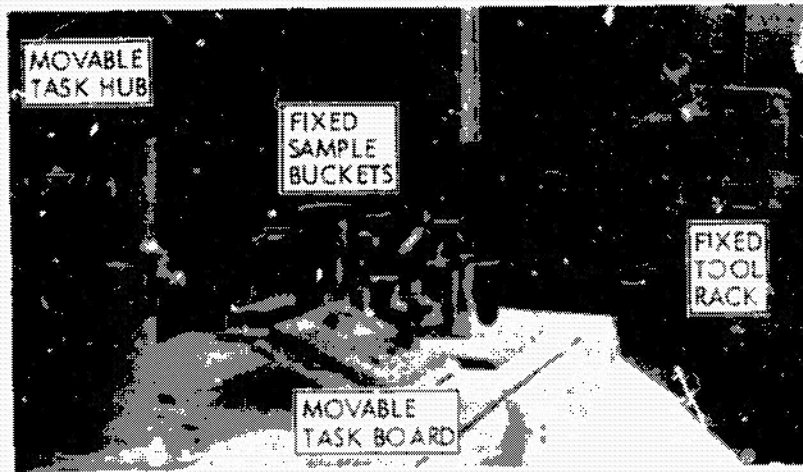

Fig. 2: Task Hub, Task Board, Sample Buckets, and Tool Rack

Experimental Design

Six basic tasks were identified for experimental investigation:[1] (1) tool retrieval; (2) tool return; (3) taking a nut off; (4) grasping an object and placing it in a container; (5) opening/closing a valve; and (6) digging. In addition, four manual control modes were delineated as important experimental parameters: (1) switch fixed rate; (2) joystick variable rate; (3) master-slave position control; and (4) master-slave position control with force reflection. With regard to the video arrangement, both mono and 2-view conditions were tested for comparison. Due to time constraints only three subjects were used for four of the tasks (tool retrieval, tool return, nut-off and sampler), and only one subject was used for the remaining two (open/close valve and digger). Each experiment was performed 5 times by each subject to obtain a statistical mean and standard deviation. Both manual and supervisory control were used.

These conditions result in a total of 1120 experimental runs. Since this would require an inordinate amount of time, the experimental load was reduced to 680 runs by noting that some of the tasks, or portions of the tasks, had constant computer execution times (see ref. 3 for details).

Subjects and Training

Three classes of subjects were used for these experiments, one experienced, four well trained, and two untrained subjects.

The well trained subjects had an average of 20 hours training given in 15 minute intervals for each of the control modes. Generally, after the subjects practiced for 15 minutes with a particular control mode a simulated task was performed. When the subjects appeared to show a plateau, experiments were begun. Since the experiments usually stretched over a period of several days, the subjects were asked to "reperform" some of the tasks due to a "mistake". If the subjects showed marked improvement the tasks were performed again until the learning curve levelled off. The four trained subjects were given incentives to perform well in the form of bonuses which would be awarded to the best combined time and error rates in any control category.

The first author was used as the baseline experienced subject. With over 200 hours of practice on manipulator systems and intimate knowledge of the SUPERMAN system, it may be reasonably assumed that the experienced subject underwent little or no learning. The experienced subject performed all of the tasks without a "warm-up" period.

The untrained subjects had a total of 3 hours training time for all control modes (i.e., 30 minutes per control mode and viewing condition). The learning curves of the untrained subjects were not observed. The only requirement placed on their training sessions was to insure that each control mode was given equal training time. After the 3 one-hour familiarity and

---

[1]Although it may appear that the tool retrieval and return tasks are simply the reverse procedure of one another, these tasks do have fundamentally different requirements. To clarify, consider that the retrieval task required the subjects to locate a 7/8 x 3/4 inch tool handle with the end effector docking plate while the return task required the subjects to mate two 1/8 inch pins and holes.

adjustment periods were over the subjects were allowed 24 hours of rest and the experiments were begun.

Procedure

The experiments were scored on the basis of recorded time and errors. The subjects were not given specific instructions to minimize either quality, but only to weigh them equally. Each subject was, however, given a criterion by which successful completion of the task would be measured (these criteria will be specified on the following pages). The experiments were not redone when errors occurred, regardless of the magnitude, unless it was impossible to proceed with the task (e.g., a collision with an object that blew a ruse, etc.). The tasks were randomized whenever possible to negate the effects of variables which the experimenters were not aware of (e.g., particularly easy or difficult task positions, short term learning effects, etc.). All tasks started from a prespecified position so that comparisons of supervisory initialization times could be made across control modes.

The procedure for each of the representative tasks was as follows:

a) **Tool-Retrieval Task** - The first task required the subject to start with the end effector positioned near the task hub. On the experimenter's signal, the subject moved the end effector to the tool rack, obtained the tool, being sure it was properly seated in the hand, and returned with the tool to the starting postion. The subjects were told that the success or failure of the task was measured by whether a solid connection between the tool handle and end effector was achieved. Execution of this task under supervisory control simply involved a button push.

b) **Tool-Return Task** - For the second task the subject started from a position next to the task hub with the tool in hand, and on the experimenter's signal, moved to the rack, replaced the tool insuring that it was properly seated, and returned to the initial position. The operators were told that the success or failure of the task was determined by whether or not the tool was properly replaced on the rack. To properly seat the tool on the rack required that both of the 1/8 inch rack pins were engaged in the handle and that the tool was completely pushed onto the pins. This task was executed under supervisory control through a simple button push.

c) **Nut-Removal Task** - This experiment began with the end effector positioned over the valve on the task hub. On the experimenter's signal, the subject moved the end effector from the valve to the nut, oriented the hand, and removed the nut. The general procedure used by the subjects and computer was to turn 180°, pull back to test if the nut was off, and then either reverse 180° and continue, or remove the nut. Prior to the task, the operators were told that the task would be considered successfully completed if the nut could be removed without losing it. Under supervisory control the operator initialized the task by moving from the starting position to the nut, orienting the hand with the rotational axis of the nut, and signaling the computer to remove it.

d) **Sampling Task** - The fourth task required the subject to pick-up thirteen randomly placed samples and put them in one of two buckets according to their size. The subjects were told that their success

or failure to complete the task would be measured by how many samples were successfully placed in the proper buckets. Under supervisory control the operator initialized the task by placing the end effector over the sample and signaling the computer to place it in the appropriate bucket. The computer returned control to the subject at the location were the sample was grasped. The operator then moved to another sample, initialized, and continued until all 13 samples were in the buckets.

e) Open/Close Valve Test - This experiment required the subject to position the end effector over the nut on the task hub, and then, on the experimenter's signal, the subject moved to the valve, oriented the hand, and opened or closed the valve as required (opening and closing tasks were switched after each experiment). The subject was required to continue until the valve operation was complete. To initialize this task under supervisory control the operator oriented the end effector on the rotational axis of the valve and signaled the computer either to open or close it as required. The computer checked the rotational torques to determine if the task had been completed.

f) Digging Task - The final task required the subject to remove a specified amount cf soil from a box by filling a bucket with a shovel. This task is composed of a number of subtasks: (1) the shovel is positioned to remove the soil, (2) the shovel is pushed into the soil and lifted out, and (3) the soil is transported to the bucket and dropped in. The subject was required to continue until the bucket was filled. Under supervisory control the positioning of the shovel was performed manually (i.e., the operator decided when and where to dig) while the scooping and dropping actions were executed by the computer.

3. RESULTS

It has been shown by a number of investigators that the time required to perform a task can be attributed to a number of distinctly different motions. For example, one classification divides the task time for control with a time delay into segments related to get, transport, and position motions [4]. For a peg-in-the-hole task Hill [5] has shown that there are two independent motions which determine the total task time under manual control - gross travel and precision. This paper will use a similar scheme to describe the task completion time for a supervisory system:

$$t_{TT} = t_I + t_P$$

where

t_{TT} = Task Time

t_I = Time required by the human operator to initialize the task. This time is primarily a function of the initial hand/task locations and the manual control mode used to locate the task.

t_P = Time required by the computer to perform the task. This time is primarily a function of the task complexity.

The determination of these times is rather simple due to the discontinuity in control which occurs during the trade from manual initialization to computer

execution (this "discontinuity" is a desired result since trading of control should be "apparent" [3,6]).

Figures 3-6 are plots of typical data (see legend below for figure abbreviations). The data recorded during the supervisory experiments have been divided into initialization and performance times to indicate the time spent by each action. Each of the time bars is the result of data averaged over two trained subjects, except for fig. 6 which is averaged over three trained subjects. The lines to the left of the manual control bars give the range over which the trained subjects performed the task. For comparison, the average time for an inexperienced subject to perform the first three tasks is also given (denoted by triangles). The mean times of the untrained subjects were always above the maximum value of the trained subjects for the same task and control mode. The lower portion of each figure (figs. 3b-6b) plots the mean number of errors which occurred under manual and supervisory control (for a specific breakdown of the individual errors see ref. 3).

LEGEND: Key to Abbreviations Used in Text

MS - Master-slave with force feedback
MS NO FFB - Master-slave without force feedback
JVRC - Joystick variable rate control
SVRC - Switch fixed/variable rate control

4. EVALUATION

Manual Control

Predictably, the task completion time increased with control complexity for all tasks. Viewing conditions (mono and 2-view) appeared to affect tasks which required precision movements (e.g., return tool and nut-off), but had little or no effect on the less precise tasks (e.g., sampling). In general, the number of errors increased as the control complexity increased from master-slave to switch rate. However, for some of the tasks a sharp decrease in errors was noticed between joystick and switch rate control (e.g., see figs. 5b and 6b). This effect is attributable to two factors: (1) the increased attention and care each operator exhibited during switch rate control modes (i.e, to move from point A to point B requires considerable thought and effort with switch rate control, but under joystick control the desired movement only requires a push on the stick), and (2) the coincidental matching of the task degrees of freedom and control degrees of freedom (e.g., in the valve or nut-off tasks the axis of rotation corresponded with the hand axis of rotation).

Table 1 gives the ratio of task completion times for each control mode with respect to the "best" control case, master-slave with force feedback. The ratios are given for each subject, task and viewing condition. The untrained subjects are denoted by U1 and U2, the trained subjects are denoted by T1, T2, T3 and T4, and the experienced subject is denoted by E1. The table shows a number of interesting trends: (1) the ratios increase with increasing control complexity, (2) the ratios are approximately constant across subjects (both trained and untrained) within a given task, (3) the ratios are constant across viewing conditions, and (4) the ratios are not constant across tasks (the tasks have been arranged in the table so that the ratio increases as the page is read from top to bottom). A number of other investigators have found similar trends [7,8,9,10].
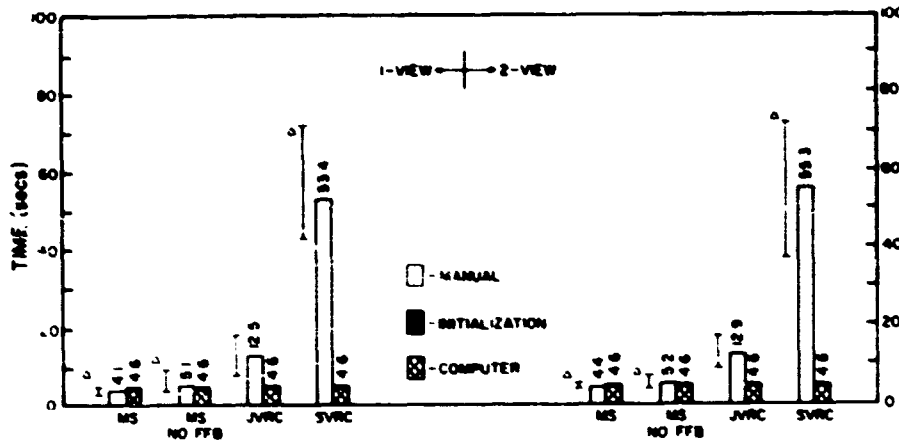
Fig. 3a: Average Tool-Retrieval Time. Each bar gives the average time
of two subjects. The Δ symbol represents the mean time for
an untrained subject. The capped lines show the total range
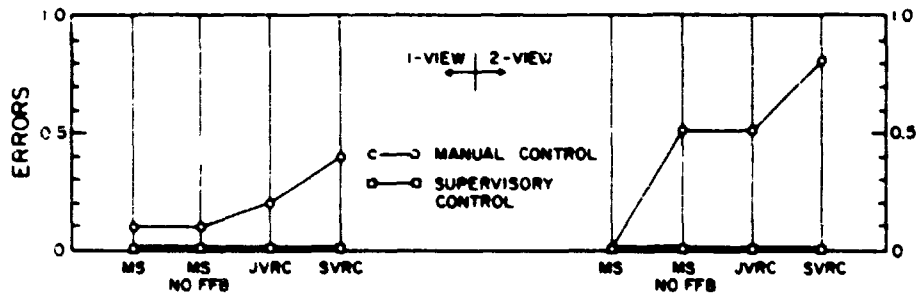of data for the trained subjects.

Fig. 3b: Expected Number of Tool-Retrieval Errors. Each data point
represents the average error rate of two trained subjects.
Possible errors included collisions, dropping the tool,
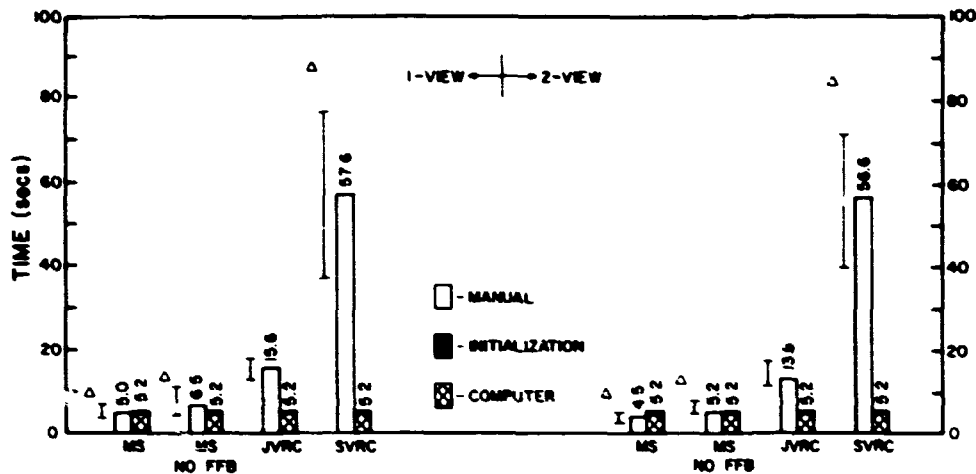and not seating the handle in the end effector properly.

Fig. 4a: Average Tool-Return Time. Each bar represents the average
time of two trained subjects and each Δ gives the mean
time for an untrained subject. The capped lines represent
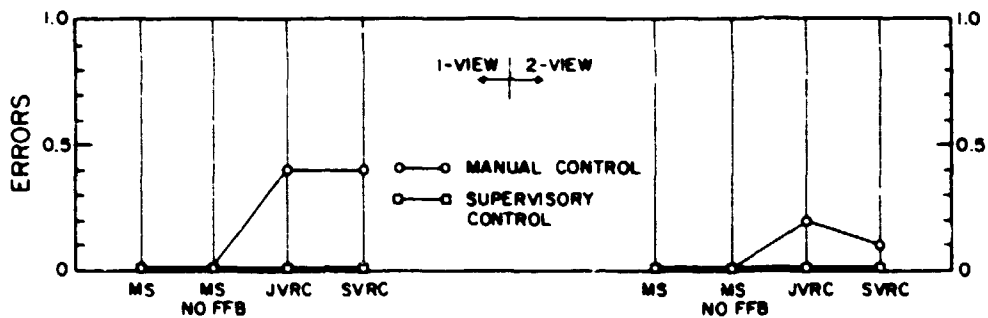the total range of data for the trained subjects.

Fig. 4b: Expected Number of Tool-Return Errors. Each data point
represents the average error rate of two trained subjects.
Possible errors included collisions, dropping the tool,
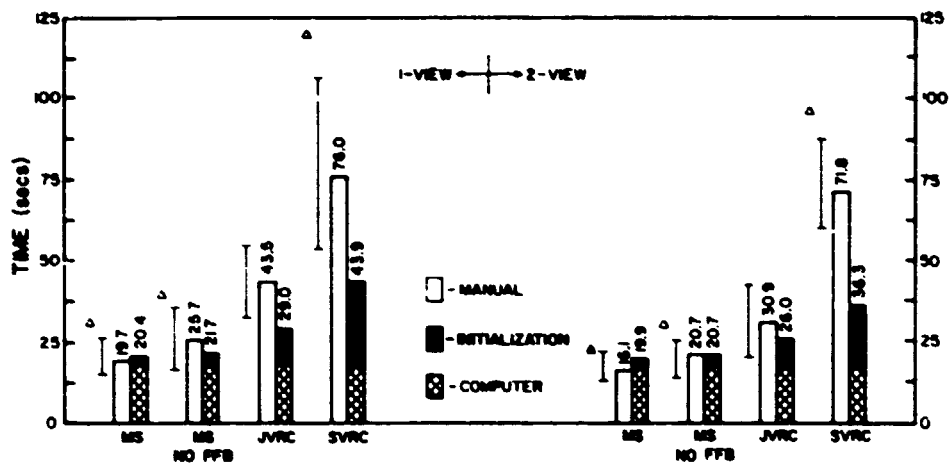and not seating the handle on the rack properly.

600

Fig. 5a: Average Nut-Removal Time. Each bar represents the average
time of two trained subjects and each Δ gives the mean
time for an untrained subject. The capped lines represent
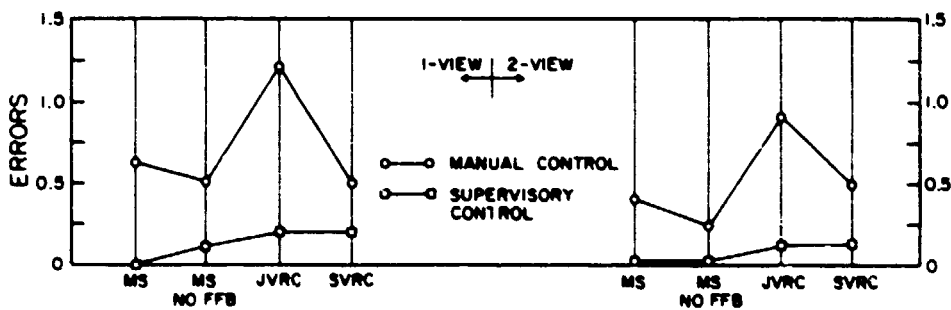the total range of data for the trained subject.

Fig. 5b: Expected Number of Nut-Removal Errors. Each data point
represents the average error rate of two trained
subjects. Possible errors included collisions and
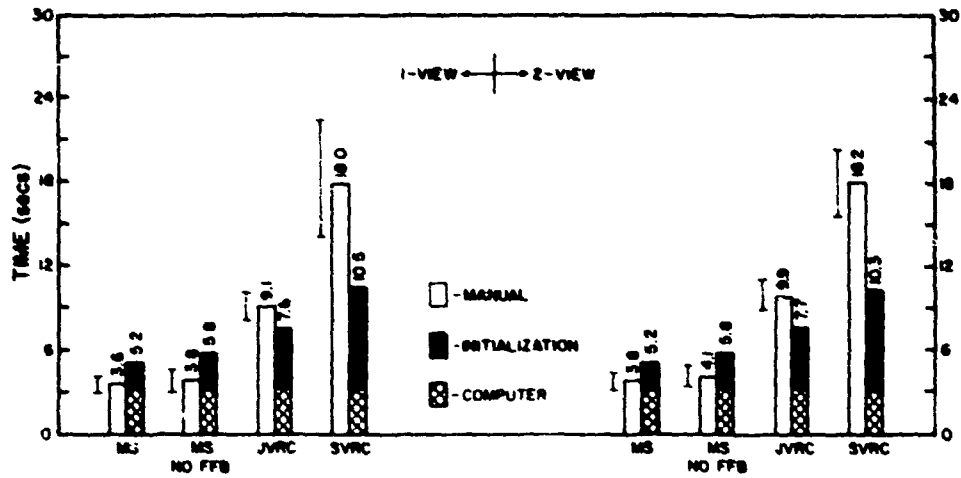dropping the nut.

Fig. 6a: Average Sampling Time. Each bar represents the mean
time of three trained subjects. The capped lines
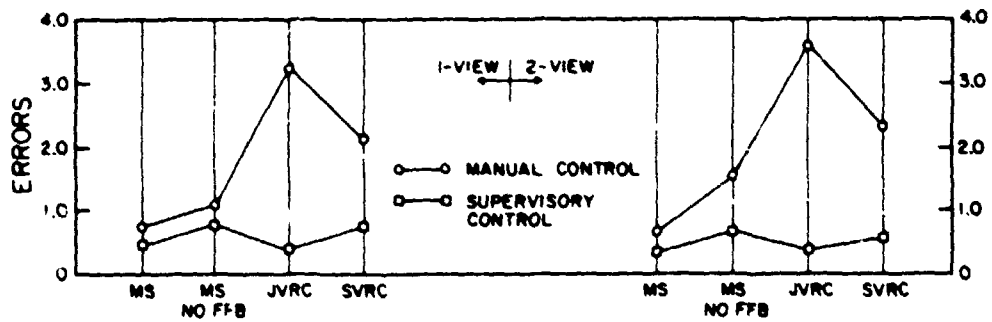represent the total range of data for the subjects.

Fig. 6b: Expected Number of Sampling Errors. Each data point
represents the average error rate of three trained
subjects for 13 sampling actions. Possible errors
included collisions, missed buckets, lost samples,
and (under supervisory control) pressing the wrong
button.

Table 1: Ratio of Time to Perform Task Under Given Control Mode to Time
to Perform Task Under Master-Slave with Force Feedback (CM/MS).

	1-VIEW				2-VIEW				
	MS	MS NO FFB	JVRC	SVRC	MS	MS NO FFB	JVRC	SVRC	
VALVE 1-DOF	1.0	1.2	1.4	3.6	1.0	1.1	1.6	3.9	E1
CUT-OFF 2-DOF	1.0	1.2	2.2	4.3	1.0	1.3	2.0	4.4	T1
	1.0	1.4	2.2	3.3	1.0	1.2	1.9	4.5	E1
	1.0	1.3	*	3.9	1.0	1.4	*	4.3	U2
SAMPLER 3-DOF	1.0	1.1	2.5	5.1	1.0	1.1	2.7	4.4	T2
	1.0	1.0	2.8	5.3	1.0	1.0	2.5	5.1	T3
	1.0	1.0	2.4	4.7	1.0	1.1	2.6	4.9	E1
SCOOPER 3-DOF	1.0	1.1	3.9	10.6	1.0	1.2	3.9	10.9	E1
RETURN-TOOL 6-DOF	1.0	1.1	3.6	10.7	1.0	1.2	3.4	12.1	T4
	1.0	1.4	2.7	12.2	1.0	1.3	2.7	13.2	E1
	1.0	1.4	*	9.0	1.0	1.4	*	9.7	U1
GET-TOOL 6-DOF	1.0	1.1	3.8	13.0	1.0	1.2	3.3	14.1	T4
	1.0	1.4	2.4	13.0	1.0	1.2	2.6	13.2	E1
	1.0	1.4	*	8.3	1.0	1.3	*	11.8	U1

Supervisory Control

As would be expected, the time required by the computer to perform its
portion of the task remained fixed regardless of the manual control mode from
which the human operator issued the execution command. Also, since the only
action required of the operator to initiate the tool-retrieval and return
tasks was a button push, the absence of initialization times in figs. 3a and
4a was not surprising. The remaining tasks, including those not shown in
this paper, had initialization times associated with the overall task time.
As seen in figs. 5 and 6 the initialization times increased with control
complexity.

Table 2 gives the ratios of the task completion times under manual
control to the times under supervisory control. The ratios are given for
each subject, task and viewing condition. The ratios relative to computer
control (Table 2) do not show the same trends as those relative to master-
slave control (Table 1). It is interesting to note that in contrast to the
consistent ratios of Table 1, the computer control ratios of the untrained
subjects are significantly higher than the trained subjects: clearly,
untrained subjects gain more from supervisory control than trained subjects.
Gains from supervisory control for any manual mode are seen to be most signi-
ficant for tasks which do not require initialization procedures other than a
button push (i.e., tool-retrieval and tool-return). The control mode columns
clearly indicate the results of the SUPERMAN experiments: (1) master-slave
with force feedback rarely benefits from supervisory control, (2) master-
slave without force feedback can profit from supervisory control in tasks
which require force feedback, and (3) both forms of rate control can be aided
by supervisory routines regardless of the task.

603

Table 2: Ratio of Time to Perform Task Under Manual Control to Time to Perform Task Under Supervisory Control (MC/SC).

	1-VIEW				2-VIEW				
	HS	HS NO FFB	JVRC	SVRC	HS	HS NO FFB	JVRC	SVRC	
VALVE 1-DOF	1.1	1.3	1 2	2 0	0.9	0.9	1 1	2.0	E1
SHUT-OFF 2-DOF	1.0 0.9 1.5	1.2 1.2 1.9	1.6 1.4 *	2.0 1.5 2.1	0.8 0.8 1.1	1 1 0.9 1.5	1.2 1.1 *	2.1 1.9 2.0	T1 E1 U2
SAMPLER 3-DOF	0.7 0.7 0 6	0.7 0.7 0 6	1 2 1 2 1 2	1.8 1.6 1 7	0.8 0.7 0.7	0.8 0.7 0.7	1.3 1.3 1.3	1.9 1.7 1.8	T2 T3 E1
SCOPPER ?-DOF	0 5	0.6	1.9	4.8	0.5	0.6	1.9	4.9	E1
RETURN-TOOL 6-DOF	0.9 1.1 1.9	1.0 1.5 2.6	3.2 2.8 *	9 4 12.8 16.9	0.9 0.9 1 7	1.0 1 1 2.4	3.0 2.3 *	10.6 11.3 16.3	T4 E1 U1
GET-TOOL 6-DOF	0.8 1.0 1.8	0.9 1.3 2.6	3.2 2.3 *	10.8 12.5 15.3	1.0 0.9 1.4	1.2 1.1 1.7	3.2 2.4 *	11.7 12.4 16 3	T4 E1 U1

In all cases the error rates for supervisory control were less than manual control. However, an interesting error was noted during the sampling experiments - occassionally the subjects pressed an incorrect button sending the sample to the wrong bucket.

5. DISCUSSION

Theoretically there is no reason why master-slave with force feedback should be any faster than supervisory control. Consider that the computer could simply mimic the human operator's best trajectory, and hence, be at least as fast. Unfortunately, in practice there is always a certain overhead associated with retransformation of coordinates, trajectory calculations and sensor logic. Also, it was generally observed that the subjects were making adaptive, orchestrated motions, whereas the computer was limited to more rigidly defined trajectories and states. In light of these observations it can be said that the faster master-slave times make more of a statement about the direction that future studies dealing with supervisory control should take than they do about its potential in teleoperator systems.

Although the experiments were not designed to measure the effectiveness of supervisory control during extended periods of manipulation, an interesting observation was made after the experiments had been completed - the manual experiments had been performed with rest periods between each run because the subjects complained of fatigue and boredom, while the supervisory experiments had been unintentionally run back-to-back since fatigue and boredom were not noted. From thes· observations it could be surmised that as a task

604

becomes more involved and complex, boredom and fatigue will become increasingly important factors, tipping the scales even further in favor of supervisory control. However, experiments to validate this statement have yet to be performed.

6. CONCLUSION

Even under "ideal" control conditions supervisory control was found to be more efficient and effective (as determined from the task completion times and manipulation errors) than switch rate control, joystick rate control, and master-slave position control. Bilateral force-reflecting master-slave was found to be slightly faster than supervisory control, but more prone to errors. Since the experiments were performed under "ideal" conditions, it can be reasonably predicted that supervisory control will show even more advantage when used with degraded sensor or control loops (e.g., time delays, limited bandwidth, etc.), though the latter experiments remain to be done. In addition, an a posteriori observation of the experimental procedure appears to indicate that the effects of operator fatigue and boredom during extended periods of manipulation can be significantly reduced through supervisory control.

References

1. Ferrel, W.R., and Sheridan, T.B., "Supervisory Control of Remote Manipulation," IEEE Spectrum, Vol. 4, No. 10, October 1967, pp. 81-88.

2. Brooks, T.L., "Supervisory Manipulation Based on the Concepts of Absolute vs Relative and Fixed vs Moving Tasks," to appear in the Proceedings of ASME International Computer Technology Conference, San Francisco, California, August 12-14, 1980.

3. Brooks, T.L., and Sheridan, T.B., "SUPERMAN: A System for Supervisory Manipulation and the Study of Human/Computer Interactions," MIT Sea Grant Report No. MITSG 79-20, July 1979, Massachusetts Institute of Technology.

4. Black, J.H., "Factoral Study of Remote Manipulation with Transmission Time Delay," Master's Thesis, M.I.T., December 1970.

5. Hill, J.W., "Two Measures of Performance in a Peg-in-Hole Manipulation Task with Force Feedback," Thirteenth Annual Conference on Manual Control, Cambridge, MA, June 1977.

6. Sheridan, T.B. and Verplank, W.L., "Human and Computer Control of Undersea Teleoperators," Man-Machine Systems Laboratory Report, Massachusetts Institute of Technology, July 1978.

7. Pesch, A.J., Hill, R.G., and Klepser, W.F., "Performance Comparisons of Scuba Divers vs. Submersible Manipulator Controllers in Undersea Work," Offshore Technology Conference, Houston, TX., 1971.

8. Mullen, D.P., "An Evaluation of Resolved Motion Rate Control for Remote Manipulators," Master's Thesis, M.I.T., January 1973.

9. Pesch, A.J., Bertsche, W.R., "Performance Measurement for Undersea Systems," Performance Evaluation of Programmable Robots and Manipulation, NBS Special Publication #459, October 1976.

10. Freedman, L.A., Crooks, W.H., and Coan, P.P., "TV Requirements for Manipulation in Space," Mechanism and Machine Theory, 1977, Vol. 12, pp. 425-438.