Proceedings of the 2003 IEEE International Workshop on
Robot and Human Interactive Communication
Millbrae, California, USA, Oct. 31 - Nov. 2, 2003

Skill Transfer in a Simulated Underactuated Dynamic Task

Marcia K. O'Malley and Abhishek Gupta

Mechanical Engineering and Materials Science
Rice University, Houston, Texas 77005, USA.
Email: omalleym@rice.edu, abhi@rice.edu

Abstract

Machine-mediated teaching of dynamic task completion is typically implemented with passive intervention via virtual fixtures or active assist by means of record and replay strategies. During interaction with a real dynamic system however, the user relies on both visual and haptic feedback in order to elicit desired motions. This work investigates skill transfer from assisted to unassisted modes for a Fitts' type targeting task with an underactuated dynamic system. Performance, in terms of between target tap times, is measured during an unassisted baseline session and during various types of assisted training sessions. It is hypothesized that passive and active assist modes that are implemented during training of a dynamic task could improve skill transfer to a real environment or unassisted simulation of the task. Results indicate that transfer of skill is slight but significant for the assisted training modes.

1 Introduction

Recently, virtual reality-based haptic training simulators have attracted the attention of a large number of researchers. The incorporation of haptic feedback in virtual reality simulations has been shown to further enhance learning rates, task performance quality, dexterity, and feelings of realism and presence over visual feedback alone [1-5]. Hence, acquisition of new motor skills might be expected to improve with the addition of haptic feedback in virtual reality based simulations. Studies, however, have shown mixed results for transfer of skill from the virtual to the real environment [5-10].

There have been relatively few studies to determine efficacy of haptic feedback in training. Yokokohji et al. investigated skill transfer in haptic training using the "What You See Is What You Feel" type of display [6]. They proposed various methods for haptic training, but found no clear evidence of skill transfer for any of the proposed methods. Adams et al. presented a study on the effectiveness of haptic training in a manual assembly task [5]. Although they found significant improvement in performance due to training, it was not clear if the improvements were attributable to the force feedback.

Several researchers have sought to use haptic devices as virtual teachers. For these applications, desired trajectories or interaction forces are programmed or recorded during an initial trial and played back to the subject in subsequent trials. Such methods rely on a two-part training procedure. First, the subject is passive while the machine guides them to complete a desired motion or displays a desired force profile. Then, the machine is passive while the subject actively tries to replicate the performance. Gillespie et al. used this record and playback method to display the optimal trajectory for balancing an inverted pendulum attached to a 1-DOF cart in the shortest possible amount of time [7]. Their results were inconclusive, though it was noted that the subjects were encouraged to adopt an alternate strategy when exposed to the teaching mode. Similar strategy was used by Kikuuwe and Yoshikawa for motion/force teaching using a fingertip presser, with limited success [8]. They indicate that subjects learned to perceive the teacher's action's clearly but not accurately.

Many haptic training studies have also focused on the display of assistive cues, often referred to as virtual fixtures, for improved human performance of positioning tasks in virtual or remote environments with force feedback. For example, Rosenberg introduced the notion of virtual fixtures as perceptual overlays for enhanced operator performance in telemanipulation tasks [11]. Virtual fixtures have also been used to improve performance in refueling tasks [12] and for manipulation assistance during interactive tasks between humans and robots [13]. Mussa-Ivaldi has shown that the application of force fields via a haptic device can cause humans to adapt to the conditions and perform desired motions [9]. He was able to show that the effects were retained for some time. Feygin et al. used haptic guidance for training users to follow complex 3D paths [10]. They found that while visual feedback was effective for teaching the trajectory shapes, haptic feedback influenced the learning of the temporal aspects of the task.

In this work, we investigate effects of haptic feedback on learning to control underactuated dynamic systems. Control of flexible structures such as ropes or springs, or of systems of masses connected by flexible links, requires the user to control more states than the number of controllable inputs. Such tasks require processing of visual and haptic data for successful completion, and are candidate tasks for rehabilitation in the instance of impaired motor control, as is the case for stroke, Parkinson's, and spinal cord injury patients. For rehabilitation applications, the end goal is direct control of the system, either virtual or real, without intervention by a therapist or embedded control system. Therefore, we present a targeting task involving control of an underactuated system of two masses connected by a spring and damper in parallel. Three assistance modes were presented to subjects during the experiments for training. These were passive assistance, shared control, and a combination of assistance and shared control. In prior work, the authors demonstrated that these assistance modes can improve human performance of a targeting task with the underactuated system [14].

This paper will present the results of an intermediate study to determine if learning of motor skills could be improved by training under various assistance modes. The transfer of skill to

0-7803-8136-X/03/$17.00©2003 IEEE

an unassisted mode after training in one of four assistance modes is presented. Human performance studies were conducted for four assistance modes while healthy subjects performed a Fitts' tapping experiment with a spring-mass-damper system [15]. Results presented in this paper present a baseline estimation of training effectiveness of each of the assistance modes.

2 Theoretical Framework

In this section, we present arguments towards the selection of an underactuated spring-mass system as a test bed for our experiments. We provide the equations of motion of the system and use the coupled dynamics of the system to derive a controller for set point control of the underactuated end-effector, which is later used in the experiments. We also provide an analysis of important aspects of the controls problem from the human operator's view point and propose an active and a passive control paradigm for faster and more accurate user control of the positioning tasks.

2.1 The Underactuated System

In real life, human beings demonstrate the capability to control complex underactuated systems with appropriate training and practice. Typical examples include pole balancing, steering control of an automobile or stable control of a yo-yo. In performance of tasks, which are essentially dynamic in nature, haptic feedback is considered to play an important role [16]. Control of a yo-yo by the user is a good example of such a case, where temporary loss of visual feedback need not affect user performance after a certain period. Robotic control of a yo-yo can only be achieved through a nonlinear controller and has been proposed as an evaluation system for intelligent controllers [17]. Given the design simplicity of the yo-yo and associated complex dynamics, a similar virtual system was implemented for our experiments.

2.1.1 System Description

The model of the underactuated system used in our study is shown in Figure 1. The system consists of two point masses connected by a flexible string modeled as a spring and a damper in parallel. The complete system was constrained to move in a plane, with no additional constraints. In the experiments described, users had complete control of the movement of one of the masses, $\mathbf{m_1}$, whereas the movement of the mass $\mathbf{m_2}$ was determined solely by system dynamics. Thus, the resulting system had 4 degrees-of-freedom (DOF) with only two of them being controlled.

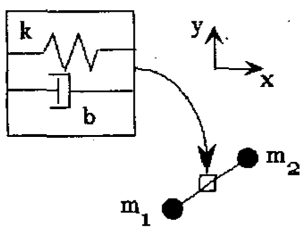

Figure 1. Virtual underactuated system. Inset shows the spring-damper model of the connection between masses m_1 and m_2.

2.1.2 Dynamics and Control of the System

The equations of motion describing the dynamics of the spring-mass system used in this study are:

$$m_1\ddot{x}_1 - F_{kx} = F_x \quad (1)$$

$$m_2\ddot{x}_2 + F_{kx} = 0 \quad (2)$$

$$m_1\ddot{y}_1 - F_{ky} = F_y \quad (3)$$

$$m_2\ddot{y}_2 + F_{ky} = 0 \quad (4)$$

where (x_1, y_1); (x_2, y_2) are positions of masses $\mathbf{m_1}$ and $\mathbf{m_2}$ respectively. Forces F_{kx} and F_{ky} denote the x and y components of the forces arising from the action of the spring and the damper; and F_x and F_y are the external forces acting on the system. In a real system, these might be the forces exerted on the mass $\mathbf{m_1}$ through actuators of an x-y table.

Without explicitly deriving the expressions for F_{kx} and F_{ky}, we now show that a simple feedback controller can be used for position control of mass $\mathbf{m_2}$. Similar to the approach proposed in [18], we define desired second order dynamics for the end effector ($\mathbf{m_2}$) as

$$\begin{aligned} \ddot{x}_2 + \lambda^2 x_2 + 2\lambda\dot{x}_2 &= -K_p(x_2 - x_d) - K_v\dot{x}_2 \\ \ddot{y}_2 + \lambda^2 y_2 + 2\lambda\dot{y}_2 &= -K_p(y_2 - y_d) - K_v\dot{y}_2 \end{aligned} \quad (5)$$

where (x_d, y_d) is the desired equilibrium point of the end effector and K_p and K_v are control gains.

Now, eliminating $\ddot{x}_2$ and $\ddot{y}_2$ from the set of equations (5) using Equations (1) through (4), we can solve for F_x and F_y. Following this analysis, our controller takes the form of Equation (6):

$$\begin{aligned} F_x &= m_1\ddot{x}_1 - m_2(2\lambda\dot{x}_2 + \lambda^2 x_2) \\ &\quad - m_2K_p(x_2 - x_d) - m_2K_v\dot{x}_2 \\ F_y &= m_1\ddot{y}_1 - m_2(2\lambda\dot{y}_2 + \lambda^2 y_2) \\ &\quad - m_2K_p(y_2 - y_d) - m_2K_v\dot{y}_2 \end{aligned} \quad (6)$$

This controller was used to verify the active assistance (shared control) mode and also the desired motion required for the subjects to successfully complete the experiments.

2.2 User Assistance

Haptic assistance to a user in performing a task can be of two types, namely passive or active. Passive assistance is characterized by use of force cues in order to convey performance information to the user as they conduct the task. This may involve use of virtual guides or force fields which exert corrective forces on the haptic interface if the operator chooses an incorrect path. This way the user is usually guided to follow the path of minimum resistance [11-12]. Active assistance on the other hand involves use of shared control or augmentation functions which help the user in actual completion of the task by applying or augmenting forces that play a role in the completion of the task [6-8, 14].

Control of an underactuated system by a human user requires learning of the dynamics of the system by the operator.

Both a passive and an active assistance strategy were designed and tested in order to study the effects of different forms of assistances. During passive assistance, a virtual force field exerted forces on the user's hand in case of a normal deviation from the desired motion of mass $\mathbf{m}_1$. The force felt was proportional to the magnitude of deviation.

During early experiments with the system described in the previous subsection, it was found that inexperienced users faced difficulty in suppressing the angular swing of mass $\mathbf{m}_2$. This was primarily due to underactuated nature of the task. Therefore, an active assistance control was implemented, based on the controller discussed earlier in the section, to suppress the swing. The active controller exerted forces on the user's hand which served to restrict deviation of mass $\mathbf{m}_2$ from the desired path (horizontal or vertical motion).

3 Methods

Fitt's Law [15] has been extensively used in literature as a measure of human performance, specifically hand-eye co-ordination. In the Fitt's law task, subjects are asked to alternately tap the centers of two rectangular objects placed side by side. The experiments are repeated with varying between-target distances and varying target sizes. The distance between the targets, the target dimensions, and the movement time are then used to characterize human performance. This task has also become a standard test for measuring performance in human-computer interaction [19-20] and in teleoperation tasks [21].

A modified Fitt's law task was used to study learning by human subjects engaged in control of the underactuated slave in the current study. The subjects were asked to repeatedly cause the end-effector, mass $\mathbf{m}_2$, to alternately hit a fixed pair of targets. More than one set of targets were used for the tests. Intuitively speaking, this should involve a rhythmic movement of the controlled mass, $\mathbf{m}_1$, similar to the control of a yo-yo. This was verified using the control system developed and presented in section 2.1.2.

3.1 Apparatus

An Impulse Engine 2000 joystick from Immersion Inc. was used to provide a high fidelity haptic simulation. The Impulse Engine has two degrees-of-freedom and a workspace of 6" x 6". The device exhibits low backdrive friction (< 0.14N) and a high sensor resolution (0.0008"). A visual display of the workspace was also created using openGL. The joystick, in conjunction with the visual feedback, was used to conduct a series of tests to study human control of underactuated systems.

An impedance control mode was used in all the experiments, as illustrated in Figure 2. Two-dimensional motion therefore was sensed at the haptic interface, which in turn displayed forces to the user. Both position and velocity were measured at the interface. It should be noted that for the simulations the joystick itself served as mass $\mathbf{m}_1$. The assistance controllers directly applied forces to the controlled mass, $\mathbf{m}_1$, which were also displayed via the haptic interface. After preliminary experiments with the authors as test subject, the values for $\mathbf{m}_2$, $\mathbf{k}$ and $\mathbf{b}$ were chosen to be 5 kg., 100 N/m and 3 Ns/m respectively, to ensure the system to be easily controllable. Forces at the haptic interface were then scaled to improve user perception. All simulations ran at the sampling frequency of 1 KHz. The system bandwidth for the apparatus is 120 Hz and it is capable of displaying a maximum force 8.9N in the workspace.

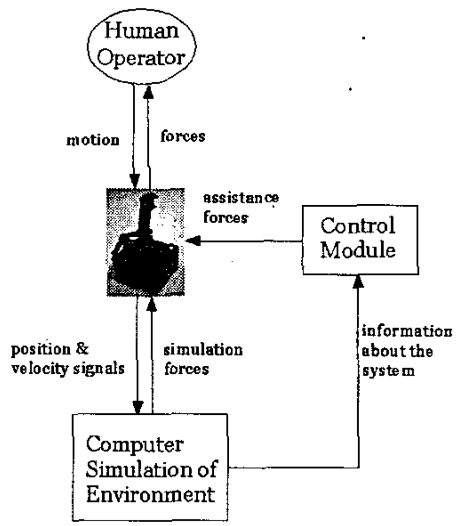

Figure 2. Block diagram of the haptic interface system. IE2000 joystick from Immersion is shown in hardware block.

3.2 Experimental Paradigms

Performance study experiments were conducted for four pairs of targets and varying forms of assistance to the user. Two target sets were aligned along the x-axis and the other two along the y-axis. Each of the targets in a set was equidistant from the origin. Thus, the subjects needed to move the joystick rhythmically, either horizontally or vertically, to alternately hit the targets in a set. There was no attempt to obstruct the subject's view of the other targets, but the active target set was marked visually.

The assistance modes chosen for the experiment were as follows:

Control Set (C): As evident from the name, this mode served as the control set. No user assistance was provided during these trials.

Passive Assistance Mode (A): In the passive assistance mode, force fields were used to encourage users to follow the horizontal or vertical direction of motion, depending upon the set of targets used. The force field generated forces proportional to either the x or y error in position of mass $\mathbf{m}_1$. These forces were subsequently displayed to the user via the joystick. The stiffness of the force field was equivalent to 22.8 N/m, with a damping of 0.57 Ns/m.

Shared Assistance Mode (S): This was the active assistance mode where a controller, similar to the one described in section 2.1.2 was implemented in simulation. The said controller applied calculated forces upon the joystick, which were in addition to the forces due to the system dynamics and were intended to prevent the swing of mass $\mathbf{m}_2$ away from the desired line of motion. For example, with the targets aligned in the horizontal direction, the controller attempts to suppress the vertical motion of mass $\mathbf{m}_2$. The inertia term corresponding to the inertia of mass $\mathbf{m}_1$ (the joystick in this case) was considered negligible, due to the high back drivability of the joystick. The following values were used for the shared control mode: $\lambda = 1$, $k_p = 70$, and $k_v = 1$.

Combined Mode (S+A): Both active and passive modes of assistance were provided to the user in this mode.

During the experiments, the subjects were asked to hit each set of targets a fixed number of times. The subjects reported that it was most important to maintain the rhythm when performing the assigned task. Human performance, under control mode settings, was recorded before and after each of the training sessions.

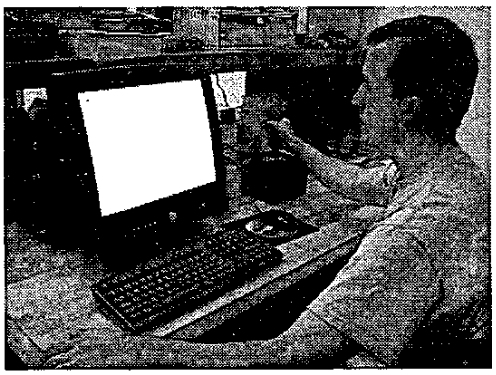

Figure 3. Test subject seated at IE2000.

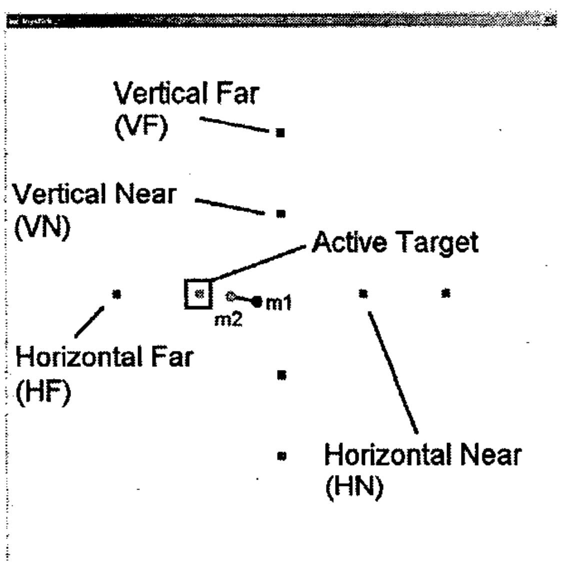

Figure 4. Graphical display of Fitts' tapping experiment with underactuated system.

3.3 Subjects

Four right-handed subjects, one female and three males, were tested. All subjects had prior experience with both the haptic interface and the task under study.

3.4 Experimental Procedure

During the course of the experiment, each subject trained with one of the assistance modes described in Section 3.2. The subject using the Control mode learned only through practice, with no passive or active assistance at any point during the trials. The experiments were conducted over a period of four days, with each subject participating at the same time each day. A single practice session was administered at the start of testing on Day 1. After the first day of testing, no further practice was allowed.

Subjects sat facing the computer screen with the dominant (right) hand holding the IE2000, as shown in Figure 3. On the first day, they were allowed to perform the tapping tasks a few times in an unassisted mode in order to become re-acquainted with the task.

A modified version of the Fitts' target tapping experiment was used to test performance, learning, and skill transfer with an underactuated system. The experiment required subjects to alternately hit a designated target pair 13 times in a row. The active target was displayed in green, while all other targets were red. Target pairs are referred to as Vertical Near (VN), Vertical Far (VF), Horizontal Near (HN), and Horizontal Far (HF) as shown in Figure 4. The assistance modes include Control (complete user control), Assistance (passive), Shared (active), and Combined (Assistance and Shared).

Each day, before and after the training session, baseline performance of the subjects was recorded by conducting the tapping test in an unassisted (Control) mode. In the baseline test, each target pair was presented once to the subject. After the initial baseline test, subjects completed two sessions of training with one of the four assistance modes. The training sessions consisted of 16 trials, with each target pair presented four times to the subject. Finally, a follow-up baseline measurement was recorded. The baseline measurements and training sessions all took place in a single sitting. For each session, the time elapsed between target hits was recorded. The order of presentation of the target pairs was randomly assigned in the computer code at the start of each baseline or training session.

4 Results

In the experiment, subjects were asked to alternately tap targets thirteen times in a row. The inter-tap times were recorded for each trial. Subjects performed an initial baseline test with one presentation of each of the four target pairs. Upon completion of the baseline test, two training sessions were administered with one of four assistance modes active. Finally, another baseline test was conducted. The before and after baseline test results for each training control mode are presented in Figures 5 through 8. Standard errors for the results are shown with error bars. Note that the baseline tests were conducted in the absence of any control mode. A two-way Analysis of Variance (ANOVA) was performed on the data presented in the figures 5 through 8.

5 Discussion

Figure 5 shows the initial and final baseline performance measurements for the Control mode, where no assistance was provided during the training sessions. This mode is equivalent to pure, unassisted practice of the dynamic target tapping task.

Results indicate that performance, measured in terms of the average inter-tap time, does not improve significantly over the four day training period (F=1.899, P=0.129), nor is improvement in performance noticeable after the training sessions on any given day (F=0.296, P=0.587).

Figures 6 through 8 present the same results for the three assisted modes. In the assistance mode, the subject's hand motion was restricted by an additional force in the direction perpendicular to the desired path (horizontal or vertical). In this mode, same day performance improvements were seen only on Day 3, but overall the improvements were not significant (F=0.182, P=0.669). Some improvements in performance were seen from Day 1 to Day 4, and differences were significant at an alpha = 0.1 level (F=2.223, P=0.075).

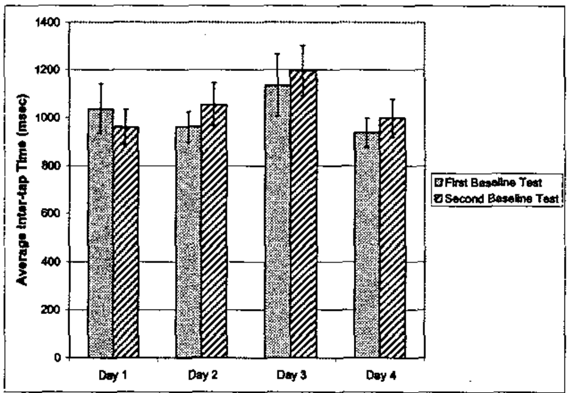

Figure 5. Control mode: Summary of baseline test performance

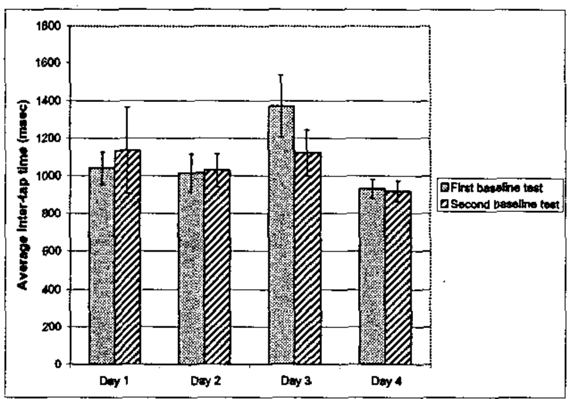

Figure 6. Assistive mode: Summary of baseline test performance

For the shared control mode, same-day performance improved consistently after each training mode, and results were significant for alpha = 0.5, as seen in the ANOVA analysis (F=3.174, P=0.025). As with the assistive control mode, slight but significant improvements were seen over the four day training period, here for alpha = 0.5 (F=3.989, P=0.047).

For the combined control mode, same day performance improved after training, with the exception of Day 4, but overall the result was not significant (F=0.941, P=0.421). However, no overall improvement in user performance was seen over the four day training period (F=1.486, P=0.224).

Although each training method was administered to a different test subject, performance levels across the control and assistance training modes were generally found to be about the same. Similarly, the shared and combined modes consistently produce average inter-tap times below one second. It is not known if this performance is due to the skill of the individual, or due to the control mode. Prior results indicate that the shared and combined modes allow for better human performance of the task, and it could be that these modes elicit better overall performance in an unassisted mode as well. Performance in general, however, improves over time for only the assistance and shared assistance modes with regard to transfer to an unassisted mode.

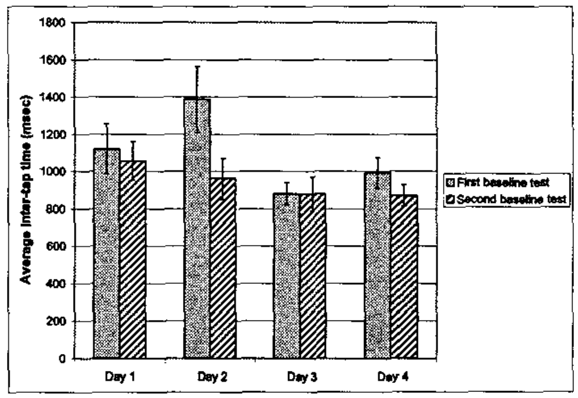

Figure 7. Shared mode: Summary of baseline test performance

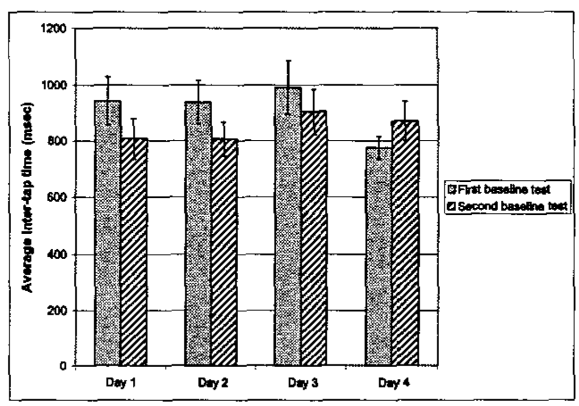

Figure 8. Combined mode: Summary of baseline test performance

The lack of improvement in performance in the combined mode could be due to a better starting performance level rather than due to a lack of training effectiveness with this control mode. In summary, there is some improvement for the shared and assistive assistance modes over the four day training period. It is suspected that the "on" or "off" approach to the assisted assistance modes, be they active or passive, is not the most effective approach to training of a dynamic task in a virtual environment. This phenomenon has been seen when augmented feedback such as computer enhancement of the environment has been added to a graphics-only virtual environment. Todorov and others noted that although some forms of augmented feedback in a graphical virtual environment were shown to enhance learning of simple movements, the performance gains achieved during training seldom transferred to the real task [22]. Future work will study the merit of adaptive assistive forces displayed in conjunction with simulated system dynamics for training or rehabilitating subjects.

6 Conclusions

We have demonstrated that both passive and active assistance during training of a dynamic targeting task serve to improve human performance of the task in an unassisted mode. Simple practice of the task without assistance did not improve performance of the task, and performance levels never reached those of the passive and active assist modes. Preliminary results indicate that active and passive assistance during training of a dynamic task may improve transfer of training over practice alone.

7 Acknowledgements

The authors wish to thank Shannon Hughes for her work on the underactuated system simulation and Atsushi Suzuki for developing the simulation graphics. We also thank the test subjects for participating in this study.

8 References

[1] Massimino, M. J. and T. B. Sheridan, "Teleoperator Performance with Varying Force and Visual Feedback." Human Factors, Vol. 36, No. 1, pp. 145-157, 1994.

[2] Meech, J. F. and A. E. Solomonides, "User Requirements When Interacting With Virtual Objects." Proceedings of the IEEE Colloquium on Virtual Reality, 1996.

[3] Fabiani, L., B. Grigore, N. Langrana, and D. Gomez, "Human Interface Using the Rutgers Master II Force Feedback Interface." Proceedings of the IEEE Virtual Reality Annual International Symposium, 1996.

[4] Richard, P. and P. Coiffet, "Human Perceptual Issues in Virtual Environments: Sensory Substitution and Information Redundancy." Proceedings of the IEEE International Workshop on Robot and Human Communication, 1995.

[5] Adams, R.J., D. Klowden, and B. Hannaford, "Virtual Training for a Manual Assembly Task." Haptics-e, Vol. 2, No. 2, http://www.haptics-e.org, October 2001.

[6] Yokokohji, Y., R. L. Hollis, T. Kanade, K. Henmi, T. Yoshikawa, "Toward Machine Mediated Training of Motor Skills – Skill Transfer from Human to Human via Virtual Environments." Proceedings of the IEEE International Workshop on Robot and Human Communication, pp. 32-27, 1996.

[7] Gillespie, R.B., M. S. O'Modhrain, P. Tang, D. Zaretzky, and C. Pham, "The virtual teacher." Proceedings of the ASME International Mechanical Engineering Conference and Exposition, Anaheim, CA, November 1998.

[8] Haptic Display Device with Fingertip Pressure for Motion/Force Teaching to Human." Proceedings of the IEEE International Conference on Robotics and Automation, pp. 868-873, 2001.

[9] Mussa-Ivaldi, F.A. and J.L. Patton, "Robots can teach people how to move arm," Proceedings of the IEEE International Conference on Robotics and Automation, pp. 300-305, 2000.

[10] Feygin, D., M. Keehner and F. Tendick, "Haptic Guidance: Experimental Evaluation of a Perceptual Motor Skill Task.," Proceedings. 10th Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems, pp. 40-47, 2002.

[11] Rosenberg, L. B., "The Use of Virtual Fixtures as Perceptual Overlays to Enhance Operator Performance in Remote Environments." Technical Report AL-TR-1992-XXX, USAF Armstrong Laboratory, WPAFB OH, 1992.

[12] Haanpaa, D.P. and G.P. Roston, "An advanced haptic system for improving man-machine interfaces." Computers and Graphics, Vol. 21, No. 4, pp. 443-449, 1997.

[13] Bettini, A., S. Lang, A. Okamura, and G. Hager, "Vision assisted control for manipulation using virtual fixtures." Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Maui, Hawaii, pp. 1171-1176, 2001.

[14] O'Malley, M. and A. Gupta, "Passive and Active Assistance for Human Performance of a Simulated Underactuated Dynamic Task," 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, pp. 348-355, 2003.

[15] Fitts, P.M., "The information capacity of the human motor system in controlling the amplitude of movement," Journal of Experimental Psychology, Vol. 47(6), pp. 381-391, June 1954.

[16] Huang, F., R.B. Gillespie, and A. Kuo, "Haptic feedback and human performance in a dynamic task," *IEEE Virtual Reality Conference*, Orlando, FL, March 24-28, 2002.

[17] Bien, Z., Y-J. Lee, S-H. Lee, K-H. Shim and S-W. Bang, "A New Benchmark System for Evaluation of Intelligent Controllers: an yo-yo System", Proceedings of 1995 IEEE International Conference on Fuzzy Systems, vol. 3, pp. 1361-1366.

[18] Reyhanoglu, M., S. Cho, and N. McClamroch, "Feedback Control of a Planar Manipulator with an Unactuated Elastically Mounted End Effector," Proceedings of the IEEE Conference on Robotics and Automation, pp. 2805-2810, 1999.

[19] Arsenault, R. and C. Ware, "Eye-hand coordination with force feedback," Proceedings of the ACM CHI, 2000.

[20] MacKenzie, I. S. , "Fitts' law as a performance model in human-computer interaction.," PhD Dissertation. University of Toronto: Toronto, Ontario, Canada, 1991.

[21] Lane, J. C., "Human Factors Optimization of Virtual Environment Attributes for a Space Telerobotic Control Station." PhD Dissertation. University of Maryland, College Park. 2000.

[22] Todorov, E., P. Shadmehr, and E. Bizzi, "Augmented feedback presented in a virtual environment accelerates learning of a difficult motor task." Journal of Motor Behavior, Vol. 29, No. 2, pp. 147-158, 1997.