Shared Control for Training in Virtual Environments: Learning Through Demonstration?

Yanfang Li* Volkan Patoglu† Marcia K. O'Malley‡

Department of Mechanical Engineering and Materials Science
Rice University
Houston, TX 77005

ABSTRACT

This paper presents an experiment to determine the underlying learning mechanism by which shared control with error reduction improves training effectiveness for dynamic manual control tasks in virtual environments. Specifically, the authors test the hypothesis that the learning mechanism of shared control with error reduction is through demonstration of the preferred strategy in the early phases of training. Three training protocols were utilized in this study and the experimental results indicate that all the benefits of shared control in training cannot be incorporated into strategy demonstration. The authors conclude that learning strategy alone is not the main mechanism of motor skill acquisition when shared control with error reduction is employed for training.

Keywords: shared control, virtual fixtures, manual control, haptic assistance, virtual training

1 INTRODUCTION

Virtual Environment (VE) technology offers a promising means of training humans for motor skill acquisition. Computationally mediated training has many potential advantages over physical training like lower risk and cost, better data collection and evaluation. Training in VE aims to transfer what is learned in the simulated environment to the equivalent real world task. Virtual training can be designed either to provide a *virtual practice medium* that matches the targeted physical medium as closely as possible, or to behave as a *virtual assistance* to improve training effectiveness by providing additional feedback in ways that are possibly not realizable in the physical world.

Most forms of interaction with computerized simulations involve only visual and auditory information. However, it is shown that the addition of haptic feedback to virtual environment simulations provide benefits over visual/auditory-only displays via reduced learning times, improved task performance quality, increased dexterity, and increased feelings of realism and presence [1–5].

To exploit training capabilities of virtual environments with haptic feedback, various *virtual assistance* paradigms have been proposed. These training paradigms are inspired by various motor learning theories and are realized through different assistance schemes such as promoting more practice, demonstrating a strategy, augmenting feedback error and reducing feedback error.

Among these methods, the most common form of haptic assist is achieved through the introduction of forbidden zones in the workspace via so called *virtual fixtures* [6]. Virtual fixtures are analogous to the use of training wheels when riding a bicycle, or a ruler when drawing straight lines. These virtual fixtures have been shown to significantly improve task performance in virtual environments [7, 8]. However, since the feedback provided by virtual fixtures is independent from the dynamics of the system to be learned, and because this feedback becomes available intermittently only to prevent large errors, from the perspective of training, virtual fixtures provide nothing more than a safer medium for practice. The assistance provided by virtual fixtures is not aimed to assist the mechanism of learning, but is designed merely to facilitate safer practice. Learning still takes place through *virtual practice*.

Another form of virtual trainer is motivated through teaching by demonstration. In these *record and play* strategies [9–12], the dynamics of an expert are recorded while performing the task and these dynamics are played backed to the novice to assist learning. In this kind of assist, the novice is not actively involved in the task during training. Once the preferred strategy to achieve the task has been played back a couple of times, the novice is allowed to practice to mimic the demonstrated dynamics. This paradigm does not account for the differences due to user-specific dynamics, and also prevents the novice from forming their own strategies.

In [13], error augmentation strategies are used to speed-up human motor learning of a dynamic task. By amplifying the instantaneous error, modified dynamics are displayed to the user to promote faster convergence of error-based adaptation mechanism. Capitalizing on a form of assistance not realizable in the physical world, this technique resulted in significant increases in learning rates. The limitation of this technique lies in its applicability to complex tasks since augmenting the error in these cases can significantly degrade performance, rendering successful task completion infeasible.

Finally, in the authors' previous work [14, 15], error reduction has been implemented through a *shared controller* for training. The authors have proposed shared control as an active assistance paradigm where the feedback is provided by a controller, which is dependent upon the system states. By dictating the type and level of active control between the computer and the human on the virtual system's dynamics, shared control constitutes the most general form of virtual training. Virtual fixtures, record and play strategies, and transient dynamics amplification are all encompassed as special cases of shared control since these paradigms can easily be realized through shared controllers of specific structures. Shared control has been shown to improve task performance in both physical and virtual environments [16, 17]. The authors' implementation of error reduction with a shared control architecture is shown to improve performance of the task as well as affecting motor skill acquisition through improved retention from one training session to the next compared to practice without assistance [15].

This paper seeks to determine the underlying learning mechanism by which shared control with error reduction improves training effectiveness for a dynamic manual control task. The authors hypothesize that a novice can benefit from an active shared controller designed to reduce the error in one of two ways: the shared controller can demonstrate the preferred strategy to successfully perform the task in the early phases of learning, or the shared controller can simplify the task dynamics resulting in a reduced number of control parameters in the adaptation mechanism, thus promoting faster learning.

*e-mail: yvonneli@rice.edu
†e-mail: vpatoglu@rice.edu
‡e-mail: omalleym@rice.edu

Facilitating learning by reducing the degrees of freedom of a complex task was first proposed by Bernstein [18]. The core idea in this hypothesis is inspired from well-developed multi-phase optimization techniques, where a coarse global search phase is followed by a fine local search that is initialized with the parameters suggested by the coarse global phase. Simplifying a complex dynamic task in early stages of training, making it significantly easier to learn (in an approximate way), and utilizing the knowledge of simplified dynamics as a useful foundation to learn the more complex task, termed *developmental progression*, has been shown to be the most effective training mechanism for neural networks [19].

In this paper, the authors test only the first hypothesis that the learning mechanism of shared control with error reduction is through demonstration of the preferred strategy in the early phases of training. Three training protocols were utilized in the study, with one group receiving no assistance throughout training (practice only), one group receiving active assistance via implementation of a shared controller that simplifies virtual system dynamics for each training session, and a third group receiving active assistance only for the first quarter of each training session. The third group is used to test the hypothesis that shared control demonstrates a preferred strategy to the subjects. Specifically, the authors investigate whether all the benefits of shared control in training can be incorporated into the strategy group which demonstrates the preferred strategy for the manual control task with a few trials in the early phase of each training session, allowing more time for practice with the uncontrolled (hence unaltered) task dynamics.

The paper is organized as follows: Section 2 describes the system and manual control task used for the experiment. The shared controller used in the experiment is introduced in Section 3. Section 4 provides details of the experimental design. The experimental results for learning of the task and statistical analysis are given in Section 5. Section 6 discusses the experimental findings. Finally, Section 7 concludes the paper.

2 System and Task Description

To determine the underlying training mechanism of shared control with error reduction, a second order manual control task of a dynamic system modeled as two point masses connected by a spring and a damper in parallel is used. This two-mass system has four degrees of freedom (DOF), namely the x and y motion of both masses m_1 and m_2. However, subjects can only control directly the x and y movement of mass m_1 via a force feedback joystick. The resulting x and y motion of m_2 is displayed graphically to the user, and is determined solely by the system dynamics. Thus, this system is an underactuated system, since the control inputs are the x and y motion of m_1.

This task is well-suited for experimental studies of human performance enhancement and training with haptic assistance because the exhibited dynamics are sufficiently complex to control but not too complex to analyze. Moreover, the force feedback generated by the interactions of the two masses connected by the spring-damper is significant for subjects to accurately control motion of the system. Haptic feedback has been shown to be an important factor for enhancing performance and learning of dynamic control tasks [20].

In this paper, besides the forces of interaction due to the system's inherent dynamics, we will also examine the effect of additional forces that we overlay on the environment for assistance due to the shared controller. Table 1 lists the three sets of system parameters that govern the dynamic response of this system. These parameter sets were varied randomly during the experiment.

Table 1: Parameters of the two mass spring damper system

Parameter Set	m_1 [kg]	m_2 [kg]	k [N/m]	b [Ns/m]
1	0	5	100	3
2	0	2	80	1
3	0	5	50	5

2.1 Hardware

An Impulse Engine 2000 joystick from Immersion Inc., shown in Figure 1, was used as the haptic display to provide high fidelity haptic simulations of the two-mass system. 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"). 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 of 8.9N in the workspace. The virtual environment graphics were created using OpenGL.

An impedance control mode was employed in all experiments, such that user motion was measured via optical encoders on the Impulse Engine, and forces were computed according to the equations of motion of the system and the additional assistance force algorithms. It should be noted that the joystick itself served as mass m_1. The displayed forces were combinations of interaction forces between m_1 and m_2 and controller assistance forces. These forces were then scaled to improve user perception.

When rendering the dynamic virtual environment, the authors neglect the inherent dynamics of the haptic device itself. This is based on two primary assumptions. First, the authors assume the device to be pseudostatic. That is, it is assumed that the motion of the haptic joystick is sufficiently slow to neglect inertial effects and Coriolis effects that are proportional to higher-order terms (velocity and acceleration). Second, the device is assumed high-quality in mechanical design and construction, such that it is free of backlash, fully backdriveable, sufficiently stiff, and of relatively low inertia.

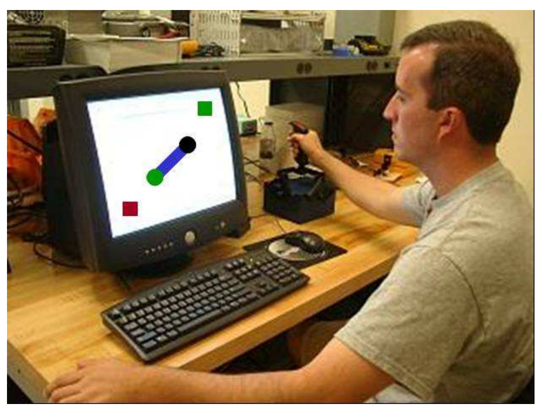

Figure 1: Subject seated at IE2000, viewing the target-hitting task.

2.2 Task

A target-hitting task is used to study manual control of the underactuated system. Subjects view the virtual environment on a computer monitor and are asked to control the motion of mass m_1 via a 2-DOF haptic joystick, thus indirectly, through the system dynamics, control mass m_2 to alternately hit a fixed pair of targets. Figure

1 shows a subject sitting in front of the haptic interface system with the virtual environment displayed on the monitor. The virtual environment display includes a pair of targets and the two-mass system. Among a target pair, one target is the active target, which is displayed in green. The other is the inactive target, displayed in red. After m_2 contacts the active (green) target, the targets change to indicate that the inactive target (red) is now active.

Figure 2 illustrates the two target pairs that are utilized in the experiments. They are referred to as follows: Positive Slope (P) and Negative Slope (N). These sloped orientations were selected because previous studies indicated that there was a significant difference in performance of the task with horizontal and vertical target orientations [14]. Each of the targets in a pair was equidistant from the origin. Therefore, the subjects needed to move the joystick (coupled to the location of m_1) rhythmically, either along the positive or negative sloped paths, to cause m_2 to alternately hit the target pair.

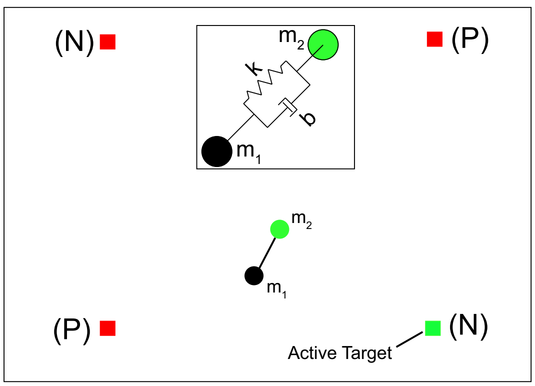

Figure 2: Graphical display of tapping experiment. Subjects control location of m_1 in order to cause m_2 to hit the desired target. Targets appear in pairs (N: negative slope; P: positive slope). Inset shows virtual underactuated system. The user controls the system by applying forces to mass m_1 through a joystick based interface.

3 Haptic Assistance - Shared Control

The goal of the experiments was to investigate the underlying mechanism of shared control with error reduction for training of the task described in Section 2. The haptic assistance is provided by additional forces displayed to the subjects via the force feedback joystick. The shared control paradigm for haptic assistance represents active intervention.

Shared control is an active assistance that depends on the dynamic system that the subject is controlling. For the task described in this paper, assistance in the form of a shared controller applies forces to the user that are a function of the desired motion of the entire virtual system and the parameters that govern the system's dynamic behavior. The shared controller implemented in this work reduces the difficulty of the task by altering the dynamics of the controlled system to help suppress the motion of the disk normal to the target axis.

Specifically, the shared controller applies forces to decrease perpendicular deviations from the preferred trajectory, forcing the motion of m2 to stay along the active target axis. Effectively, the action of the shared controller is to feed the constraint forces to be imposed on m_2 to m_1 (hence to the subject) via the inverse dynamics of the dual mass-spring-damper system. Details of implementation of this shared controller can be found in [15].

4 Experimental Details

Twelve subjects (2 female, 10 male, ages 22-28), primarily graduate students in engineering, participated in the experiments. These subjects were divided into three different groups, each group with four subjects. Group *N* is referred to as the no assistance group (control group) that receives no haptic assistance during training portion of the manual control task. Subjects in Group *S* called strategy group receive shared control as active assistance during roughly the first 25% of each training portion. Subjects in Group *A*, shared control group, receive shared control for assistance for the duration of each training portion of the session.

The hypothesis we proposed is as follows: Shared control during manual task training serves to demonstrate a preferred strategy that subjects are able to adopt. This is the underlying learning mechanism of shared control.

If this hypothesis holds, we would expect the strategy group (S) to outperform both the shared control (A) and no assistance (N) groups during baseline measurements. These baseline measurements are administered before and after each training portion to assess user performance of the task in an unassisted mode (see Figure 3). Better performance than no assistance (N) group is expected since strategy (S) group gets the benefit of demonstration of a preferred strategy for task completion. Better performance than shared control (A) group is expected since strategy (S) group has more time to interact with the baseline (unassisted) system dynamics.

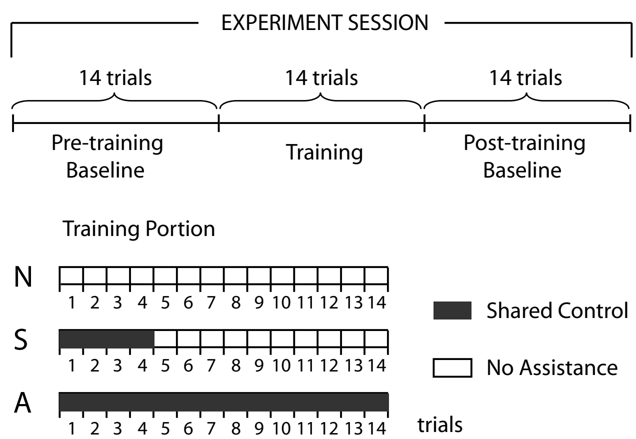

Figure 3: Schematic representation of the training sessions design for the experiment. Each training session contains three portions: pre-training baseline, training, and post-training baseline. During each training session the no assistance (N) group receives no assistance whereas the strategy (S) group is provided with shared control assistance in the first 4 trials over 14 trials of the training portion and the shared control (A) group is provided with shared control assistance throughout all 14 trials of the training portion.

The experiment includes an evaluation session, nine training sessions, and a retention session. In order to control individual differences in performance across subjects, each subject performed the task during an evaluation session, administered without haptic assistance. The purpose of the evaluation session was to group the subjects based on their initial performance of the task. Before the evaluation session, subjects are given a maximum of five minutes to become familiar with the haptic joystick and the virtual environment. After this introductory period, all subjects complete an identical evaluation session to determine initial performance. During the evaluation session, the three sets of system parameters (see Table 1) are presented randomly to the subject with ten repetitions, for a total of 30 trials, each with duration of 20 seconds. Each subject is scored based on the total number of target hits. These subjects are then ranked according to the score, then divided into quartiles by ranking. Subjects from each quartile are then randomly assigned into the three groups (no assistance (N), strategy (S), and shared control (A)) such that the average score for three groups are roughly equivalent at the start of training.

All groups completed nine training sessions, each training session containing three portions: pre-training baseline, training, and post-training baseline. Moreover, each training portion consists of 14 trials as shown in Figure 3. The no assistance (N) group serves as the control set and no haptic assistance was provided during the training portion. The "no assistance" case is akin to practice. In this interaction mode, subjects feel the forces generated solely due to the internal dynamics of the system. In contrast, for the shared control case, subjects felt the forces due to both the internal dynamics of the system and the augmented forces intended to assist in task completion during the training portion of the session. The strategy (S) group is provided with shared control assistance in the first 4 trials over 14 trials of the training portion whereas the shared control (A) group is provided with shared control assistance throughout all 14 trials of the training portion.

In order to assess the improvement of subjects across the nine training sessions, a baseline test, in which no assistance was applied, was completed before and after each training portion. For each baseline test, subjects completed 14 trials, all in no assistance mode. A training portion and its corresponding pre- and post-training baseline tests took place in a single sitting. The nine training sessions were separated by two to three days, such that subjects completed all sessions in no less than three but no more than four weeks.

One month after the final training session, all subjects completed one retention session. In the retention session, no haptic assistance was provided; subjects felt only the interaction forces between the two masses. The retention session was conducted identical to the evaluation session, with thirty trials of twenty seconds each, and each system parameter set presented ten times in random order to the subject.

5 Results

Figure 4 presents trajectories of mass m_2 at the first and the last trial of training sessions for a typical subject. Subjects adapt a preferred strategy where their trajectories converge to a straight-line path between targets. All the subjects except one subject in no assistance group show a similar trend, which implies that they adapt to the preferred strategy by the end of the training course.

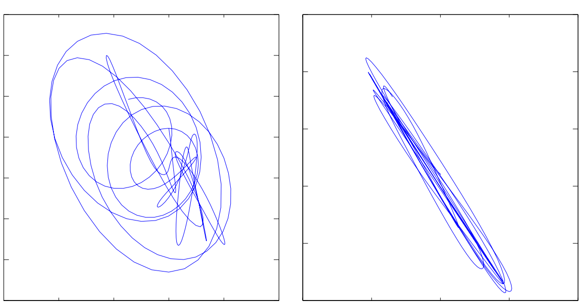

Figure 4: Trajectories of mass m_2 for a typical subject. The figure on left presents the trajectory during the first trial and the figure on the right depicts the trajectories during last trial of training sessions.

Two performance measures are used to assess subject performance of the target-hitting task. These measures are the total target hit count and average error. Specifically, hit count is the total number of target hits within one twenty-second trial. The average error is the average of the position deviation of the mass m_2 from the target axis. This average error measure is dependent on the strategy adopted by the user.

Figures 5-7 show the experiment results in terms of hit count for pre-training baseline (Fig. 5), training portion (Fig. 6), and post-training baseline (Fig. 7). The error bars indicate the standard errors for the results. In the following figures *A* stands for shared control group, *N* for no assistance group, and *S* for strategy group.

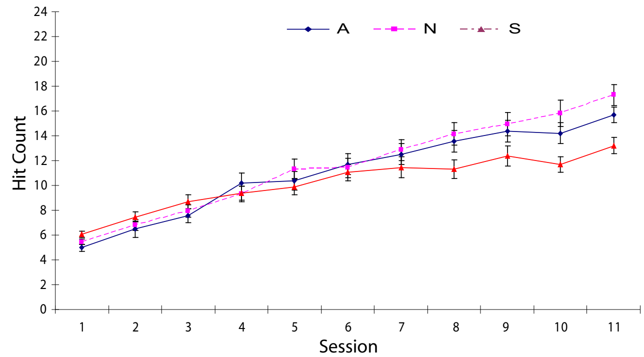

Figure 5: Pre-training baseline hit count for different groups over eleven sessions (including the evaluation session (session #1), nine training sessions (sessions #2-10) and the retention session (session #11)). Group A represents the shared control group, N is the no assistance group, and S is the strategy group, which receives shared control for assistance for the first quarter of each training session.

In Figure 5, pre-training baseline hit counts for different groups over eleven sessions including the evaluation session (session #1), nine training sessions (sessions #2-10), and the retention session (session #11) are shown. As can be seen from the evaluation session, all three groups starts at approximately the same performance level in terms of hit count. Figures 6 and 7 show the hit counts for different groups over nine training sessions for the training portion and post-training baseline, respectively. During these nine training sessions, subject performance continues to improve. The retention data in Figure 5 shows that learning continues even one month after the last training session. The learning effect at the end of the experiment is significant, starting from approximately five hits per trial and improving to sixteen hits in the retention session.

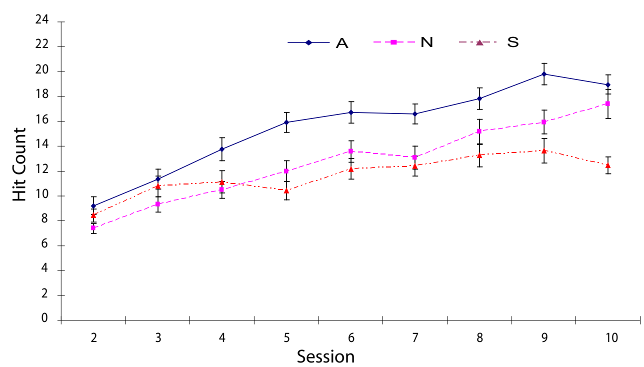

Figure 6: Training portion hit count for different groups over nine training sessions (sessions #2-10). Shared control (A) group outperforms no assistance (N) and strategy (S) groups during training.

Figures 8 to 10 are analogous to Figures 5 to 7 with performance results given in terms of the average error measurements. Since the

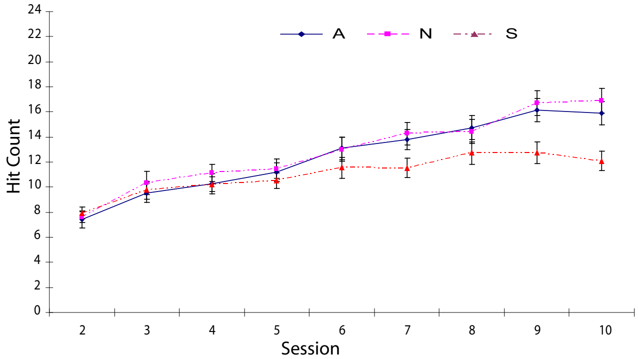

Figure 7: Post-training baseline hit count for different groups over nine training sessions (session #2-10). Shared control (A) and no assistance (N) groups outperform strategy (S) group.

score used in evaluation session to group subjects is based on the hit count instead of average error, the starting average error for the three groups is different. The no assistance group need not show a learning trend in this metric since this average error measure is dependent on the preferred strategy. However, the results suggest that most of the subjects in no assistance group improved in this metric as can be seen in Figures 8 to 10. There is one particular subject, called N_1, who failed to adapt to the preferred strategy. N_1 uses a strategy that manipulates the joystick elliptically to achieve target hitting. The performance of this subject is noticeably different from others in both measures (hit count and error). Subject N_1's hit count and average error measurement remain nearly constant throughout entire training course. Due to this alternate strategy adopted by N_1, the learning curve for the average error measure of the no assistance group N does not show any improvement and contains relatively large error bars. To remove the effects of N_1 on the average measure performance of the non-assistant group, Figures 8 to 10 include plots for the non-assistant when N_1 is ruled out, depicted as N'.

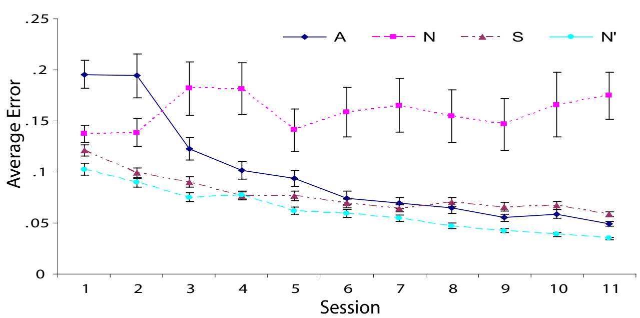

Figure 8: Pre-training baseline average error for different groups over eleven sessions (the evaluation session (session #1), nine training session (sessions #2-10), retention session (session #11)). Group A represents the shared control group, N is the no assistance group, and S is the strategy group, and N' is the no assistance group without subject N_1's data.

The baseline performance measured by both the hit count and average error shows that the shared control (A) group outperforms the strategy (S) group. When the results for subject N_1 are excluded, the shared control (A) group exhibits similar performance as the no assistance (N) group in terms of hit count and average error. Additionally, the no assistance (N) group demonstrates better performance than strategy (S) group both in hit count and average error.

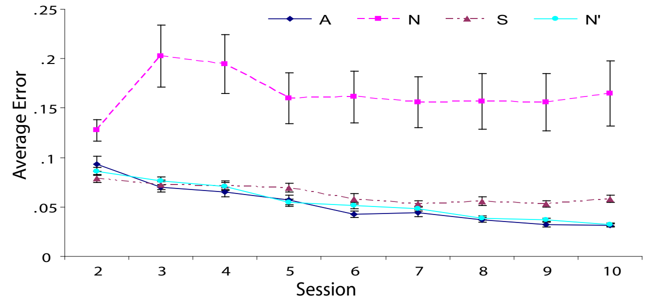

Figure 9: Training portion average error for different groups over nine training sessions (sessions #2-10). Shared control (A) group outperforms strategy (S) and no assistance without N_1 (N') groups during training.

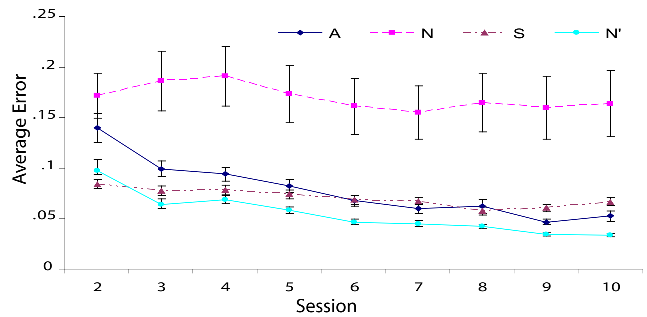

Figure 10: Post-training baseline average error for different groups over nine training sessions (sessions #2-10). Shared control (A) and no assistance without N_1 (N') groups outperform strategy (S) group.

A repeated measures Analysis of Variance (ANOVA) was carried out for the experimental results, and included F-test contrasts that are constructed simultaneously in order to evaluate a set of custom hypotheses regarding the main effects. These contrasts also control for experiment-wise error rate, whereas independent t-tests do not. Two-way interactions were analyzed, with significant interactions for the following combinations: the effect between no assistance (N) group and strategy (S) group is significant, $p = 0.0081$ and $p = 0.0241$. However, the P value based on average error indicates that although shared control (A) group and no assistance (N) group have the same hit count performance, the average error effect between these two groups are significant, $p = 0.0087$ for pre-training baseline and $p < 0.0001$ for post-training baseline.

6 DISCUSSION

The experimental results with hit count as the performance measure show that the performance of shared control (A) group and strategy (S) group are significantly different ($p = 0.0766$) for post-training baseline. Additionally, the shared control (A) group has better performance than strategy (S) group in terms of hit count. This result contradicts the hypothesis proposed in Section 4. The differences between these two groups indicate that the subjects learn this manual control task not simply through demonstration of the preferred strategy in the early training portion of the experiment.

Moreover, the statistical analysis based on average error shows that shared control (A) group and strategy (S) group are not significantly different ($p = 0.9883$) for post-training baseline, which indicates that both of these groups adapt the preferred strategy demonstrated by the shared controller. Since both of these groups adapt the preferred strategy but exhibit different performance in hit count, learning strategy alone is not necessarily the main mechanism of motor skill acquisition for this manual control task. However, demonstration of strategy might have a beneficial effect on task learning. The subject N_1 in no assistance (N) group was never exposed to the preferred strategy and never adopted it. This subject (N_1) gained no improvement in performance throughout training although improvements with alternative strategies are possible.

Furthermore, the performance of shared control (A) group is better than strategy (S) group. The only difference in assistance for these two groups lies in the percentage of exposure to shared control during the training portion. Therefore, the dose of shared control is an important factor for training. The trends in the learning curves indicate that improvement in both shared control (A) and strategy (S) groups are similar at the early stages of training. However, at a certain point in training, the learning curves start to diverge from each other reaching different performance levels at the end of training. These results suggest that there exists an optimal dose to deliver shared control. In our further studies, we will pursue an adaptive shared controller based on subject performance.

The results also demonstrate that strategy (S) group is significantly different from no assistance (N) group ($p = 0.0241$), with the no assistance (N) group outperforming the strategy (S) group. This is quite interesting, since the only difference between these two groups is that strategy group had exposure of shared control with only 4 trials over 42 trials per session. Yet the small difference in exposure of shared control results in significantly different performance. In contrast to the strategy (S) group, the shared control (A) group exhibits no significant difference with no assistance (N) group in terms of hit count. The interaction with secondary dynamics while trying to learn primary dynamics might result in interference or consolidation [21, 22]. The amount of exposure to secondary dynamics is an important factor for these interactions. Therefore, one possible explanation for poor performance of strategy (S) group might be due to the fact that active assist in the first four trials of the training session introduced interference, thereby undermining the performance in the baseline tests. The interference effect seems to be non-existent in shared control (A) group.

The performance of shared control (A) group is not significantly different from the no assistance (N) group, which might be attributed to the simplicity of the manual control task. Facilitating learning by reducing task complexity is valuable only for complex tasks [18]. A more complex task is required to emphasize the benefits of shared control with error reduction, which will be pursued by the authors in the future.

7 Conclusion

An experiment was presented to investigate the underlying learning mechanism by which shared control with error reduction improves training effectiveness for dynamic manual control tasks in virtual environments. Three training protocols were utilized to test the hypothesis that learning mechanism of shared control with error reduction is through demonstration of the preferred strategy in the early phases of training. The three protocols included practice in the haptic virtual environment (no assistance), training with haptic assistance (shared control) throughout all sessions, and training with haptic assistance for the first quarter of each training session as a means of demonstrating a preferred strategy while still allowing for unassisted practice of the task. The experimental results indicate that learning strategy alone is not the main mechanism of motor skill acquisition for this task, and all the benefits of shared control in training cannot be incorporated into strategy demonstration.

Acknowledgements

The authors wish to thank Dr. Scott Baggett of the Statistical Consulting Lab at Rice University for his support in experimental design and analysis of results. This work was supported in part by the Office of Naval Research Grant number N00014-04-1-0517.

References

[1] M. J. Massimino and T. B. Sheridan, "Teleoperator performance with varying force and visual feedback," *Human Factors*, vol. 36, no. 1, pp. 145–157, 1994.

[2] P. Richard and P. Coiffet, "Human perceptual issues in virtual environments: sensory substitution and information redundancy," in *IEEE International Workshop in Robot and Human Communication*, pp. 301–306, 1995.

[3] J. F. Meech and A. E. Solomonides, "User requirements when interacting with virtual objects," *Virtual Reality - User Issues*, vol. 68, no. 1, pp. 1–3, 1996.

[4] R. J. Adams, D. Klowden, and B. Hannaford, "Virtual traning for a manual assembly task," *Haptics-e*, vol. 2, no. 2, 2001.

[5] M. K. O'Malley, K. J. Hughes, D. F. Magruder, and R. O. Ambrose, "Simulated bilateral teleoperation of robonaut," in *AIAA Space 2003 Conference*, 2003.

[6] L. B. Rosenberg, "Virtual fixtures: perceptual tools for telerobotic manipulation," in *IEEE Annual Virtual Reality International Symposium*, pp. 76–82, 1993.

[7] D. P. Haanpaa and G. P. Roston, "An advanced haptic system for improving man-machine interfaces," *Computers & Graphics*, vol. 21, no. 4, pp. 443–449, 1997.

[8] A. Bettini, P. Marayong, S. Lang, and A. M. Okamura, "Vision assisted control for manipulation using virtual fixtures: Experiments at macro and micro scales," in *IEEE International Conference on Robotics and Automation*, pp. 3354–3361, 2002.

[9] R. B. Gillespie, P. T. S. O'Modhrain, C. Pham, and D. Zaretsky, "Virtual teacher," in *ASME International Mechanical Engineering Congress and Exposition*, pp. 3354–3361, 1998.

[10] Y. Yokokohji, R. Hollis, T. Kanade, and K. Henmi, "Toward machine mediated training of motor skill," in *IEEE International Workshop on Robot and Human Communication*, pp. 32–37, 1996.

[11] K. Henmi and T. Yoshikawa., "Virtual lesson and its application to virtual calligraphy system," in *IEEE International Conference on Robotics and Automation*, pp. 1275–1280, 1998.

[12] R. Kikuuwe and T. Yoshikawa, "Haptic display device with fingertip presser for motion/force teaching to human," in *IEEE International Conference on Robotics and Automation*, pp. 868–873, 2001.

[13] P. S. E. Todorov 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.

[14] M. K. O'Malley and A. Gupta, "Passive and active assistance for human performance of a simulated underactuated dynamic task," in *IEEE International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 348–355, 2003.

[15] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared control in haptic systems for performance enhancement and training," *ASME Journal of Dynamic Systems, Measurement and Control (JDSMC) Special Issue on Novel Robotics and Control*, vol. 128, no. 1, pp. 75–85, 2006.

[16] P. Griffiths and R. B. Gillespie, "Sharing control between human and automation using haptic interface: Primary and secondary task performance benefits," *Human Factors*, vol. 47, no. 3, pp. 574–590, 2005.

[17] M. Yoneda, F. Arai, T. Fukuda, K. Miyata, and T. Naito, "Assistance system for crane operation with haptic display: Operational assistance to suppress round payload swing," in *IEEE International Conference on Robotics and Automation*, pp. 2924–2929, 1999.

[18] N. Bernstein, *The Coordnation and Regulation of Movements*. Oxford: Pergamon, 1967.

[19] V. Ivanchenko and R. A. Jacobs, "A developmental approach aids motor learning," *Neural Computation*, vol. 15, pp. 2051–2065, 2003.

[20] F. Huang and R. B. Gillespie, "Visual and haptic feedback enable online tuning to resonant dynamics," *Journal of Motor Behavior*, in review.

[21] R. Shadmehr, T. Brashers-Krug, and F. Mussa-Ivaldi, "Interference in learning internal models of inverse dynamics in humans," *Advance in Neutral Information Processing Systems*, pp. 1117–1124, 1995.

[22] T. Brashers-Krug, R. Shadmehr, and E. Bizzi, "Consolidation in human motor memory," *Nature*, pp. 252–255, 1996.