Negative Efficacy of Fixed Gain Error Reducing Shared Control for Training in Virtual Environments

YANFANG LI
Rice University
VOLKAN PATOGLU
Sabancı University
and
MARCIA K. O'MALLEY
Rice University

Virtual reality with haptic feedback provides a safe and versatile practice medium for many manual control tasks. Haptic guidance has been shown to improve performance of manual control tasks in virtual environments; however, the efficacy of haptic guidance for training in virtual environments has not been studied conclusively. This article presents experimental results that show negative efficacy of haptic guidance during training in virtual environments. The haptic guidance in this study is a fixed-gain error-reducing shared controller, with the control effort overlaid on the dynamics of the manual control task during training. Performance of the target-hitting manual control task in the absence of guidance is compared for three training protocols. One protocol contained no haptic guidance and represented virtual practice. Two protocols utilized haptic guidance, varying the duration of exposure to guidance during the training sessions. Exposure to the fixed-gain error-reducing shared controller had a detrimental effect on performance of the target-hitting task at the conclusion of a month-long training protocol, regardless of duration of exposure. While the shared controller was designed with knowledge of the task and an intuitive sense of the motions required to achieve good performance, the results indicate that the acquisition of motor skill is a complex phenomenon that is not aided with haptic guidance during training as implemented in this experiment.

Categories and Subject Descriptors: H.1.2 [**Model and Principles**]: User/Machine Systems—*Human factors*; H.5.2 [**Information Interfaces and Presentation**]: User Interfaces—*Haptic I/O, theory and methods, evaluation/methodology*; K.4.2 [**Computers and Society**]: Social Issues—*Assistive technologies for persons with disabilities*

General Terms: Experimentation, Human Factors

Additional Key Words and Phrases: Shared control, manual control, haptic assistance, virtual training, motor skill training

ACM Reference Format:

Li, Y., Patoglu, V., and O'Malley, M. K. 2009. Negative efficacy of fixed gain error reducing shared control for training in virtual environments. ACM Trans. Appl. Percpt. 6, 1, Article 3 (February 2009), 21 pages. DOI = 10.1145/1462055.1462058 http://doi.acm.org/10.1145/1462055.1462058

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

This work was supported in part by the Office of Naval Research, grant number N00014-04-1-0517.
Authors' addresses: Y. Li and M. K. O'Malley (corresponding author), Rice University, Houston, TX 77005; email: {yvonneli, omalleym}@rice.edu; V. Patoglu, Sabancı University, İstanbul, Turkey 34956; email: vpatoglu@sabanciuniv.edu.

like lower risk and cost and improved data collection via integrated sensors that provide a means for objective evaluation. 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 or auditory displays via reduced learning times, improved task performance quality, increased dexterity, and increased feelings of realism and presence [Massimino and Sheridan 1994; Richard and Coiffet 1995; Meech and Solomonides 1996; Adams et al. 2001; Williams et al. 2002; O'Malley et al. 2003].

Virtual training can be designed either to provide a virtual practice environment that matches the targeted physical environment as closely as possible or to provide virtual assistance intended to improve training effectiveness. Regardless of the approach, the aim of training in VEs is to transfer what is learned in the simulated environment to the equivalent real world task. Caution should be taken when using virtual environments for training, since it has been shown in the literature that intuitive training schemes in computationally mediated environments with visual and auditory feedback may not result in positive transfer effects and can even lead to negative transfer [Kozak et al. 1993; Lintern 1991; Lintern and Roscoe 1980; Gamberini 2000]. Negative transfer effects are attributed mainly to limitations in the fidelity of the virtual task compared to the real task due to simplifications required for rendering. Negative transfer effects may also be attributed to the augmentation of task dynamics due to the presence of virtual guidance. Despite the potential limitations of virtual environment training, numerous positive training outcomes have been reported as further outlined below, especially for virtual assistance schemes.

To exploit training capabilities of virtual environments with haptic feedback, various virtual guidance paradigms have been proposed. These training paradigms are inspired by several 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 guidance is achieved through the introduction of perceptual constraints on the workspace via so called virtual fixtures [Rosenburg 1993]. Virtual fixtures have been shown to significantly improve task performance in virtual environments [Haanpaa and Roston 1997; Bettini et al. 2002]. However, virtual fixtures are not effective for training, since the user becomes dependent on the existence of virtual fixtures to complete the task [O'Malley et al. 2006]. To overcome this limitation of virtual fixtures in training, dead zones are widely implemented. Analogous to training wheels when riding a bicycle, virtual fixtures with dead zones introduce forbidden regions to the task space. Virtual fixtures with dead zones provide improved training efficacy, since the feedback provided is independent from the dynamics of the system to be learned. Additionally, due to the dead zones, haptic guidance becomes available intermittently only to prevent large errors. From the perspective of training, this kind of assistance provides nothing more than a safe medium for practice. The assistance provided by virtual fixtures with dead zones is not intended to assist the mechanism of learning, since learning still takes place through virtual practice.

Another form of virtual trainer is motivated through teaching by demonstration. In these record and play strategies [Gillespie et al. 1998; Yokokohji et al. 1996; Henmi and Yoshikawa 1998; Kikuuwe and Yoshikawa 2001; Feygin et al. 2002], the dynamics of an expert are recorded while performing the task. The dynamics are then played backed to the novice to assist learning. In this training scheme, the novice remains passive while the desired motor skills are displayed. Once the preferred task completion strategy has been displayed to the novice a sufficient number of times, the novice is asked to mimic the demonstrated dynamics. The record and replay training paradigm does not account for differences due to user-specific dynamics and also prevents the novice from forming their own strategies for task completion. Results from studies on record and replay effectiveness for motor skill training are highly inconclusive [Gillespie et al. 1998; Yokokohji et al. 1996; Henmi and Yoshikawa 1998; Kikuuwe and Yoshikawa 2001; Feygin et al. 2002].

While virtual fixture and record and replay methods may be successful for training, they do not take advantage of the ability of a haptic interface to display state dependent feedback during manual interactions. Virtual fixtures with dead zones provide intermittent penalties when the user violates the territory of the forbidden region. However, the penalty depends only on the user position and is independent of all other states of the coupled system. On the other hand, while the record and replay strategy is capable of demonstrating higher-order control schemes than the virtual fixture approach, the subject remains passive during the replay mode. When the subject is actively moving through the devices workspace, no corrective feedback is provided to the user.

The authors have proposed shared control for training as an active haptic guidance paradigm where feedback is provided by a controller, which is dependent upon the system states [O'Malley and Gupta 2003]. A shared controller dynamically intervenes, through an automatic feedback controller acting upon the system, to modify the (coupled) system dynamics during training. Shared controllers can take many different forms and can modify the system dynamics in different ways. By dictating the type and level of active control between the computer and the human on the virtual system's dynamics, training with shared control constitutes the most general form of virtual training. Virtual fixtures and record and play strategies are special cases of shared control since these paradigms can be realized through shared controllers of specific structures. Examples of shared control for training are discussed in the following paragraphs.

Shared control has been used to train reaching movements by generating custom force fields designed to drive subjects to adapt a predefined trajectory [Patton and Mussa-Ivaldi 2004]. This strategy is based on aftereffects of adaptation and aims to alter the feedforward command in the central nervous system. However, these perturbing force fields have not been shown to be effective for long-term training, since the aftereffects tend to wash out after relatively short periods.

In Todorov et al. [1997], error amplification strategies are used to speed up human motor learning of a dynamic task. Patton et al. [2001] amplified directional errors during reaching movements with a robotic device to improve motor learning after stroke. Similarly, Emken and Reinkensmeyer [2005] utilized a deadbeat transient amplification paradigm to accelerate adaptation to a novel dynamic environment. In all these paradigms, modified dynamics are displayed to the user to promote faster convergence of error-based adaptation mechanisms through amplification of the instantaneous error. The error-amplification techniques, which capitalize on a form of haptic guidance not realizable in the physical world, resulted in significant increases in learning rates. Error amplification techniques are limited in their applicability to more complex tasks, since augmentation of error could significantly degrade performance, rendering successful task completion infeasible.

Finally, in the authors' previous work [O'Malley and Gupta 2003; O'Malley et al. 2006], error reduction has been implemented through a fixed gain shared controller. Reduction of error simplifies the target-hitting task by reducing the degrees of freedom of the controlled system. The task simplification approach is advocated in Lintern [1991] as long as the important perceptual invariants of the task are preserved. In the error reduction implementation of shared control, the dynamics of the (state dependent) shared controller are designed such that the coupled (closed loop) dynamics of the system are simpler to manipulate than the system dynamics without the controller in place. Hence, by simplifying the task dynamics through coupling, the shared controller helps the user to achieve better task performance. Schematic representations of a sample error reducing shared control implementation for training of a bicycle riding task in physical and virtual environments are given in Figures 1(a, b). In Figure 1(a), the dynamics of the physical bicycle balancing task are illustrated as an inverted pendulum,

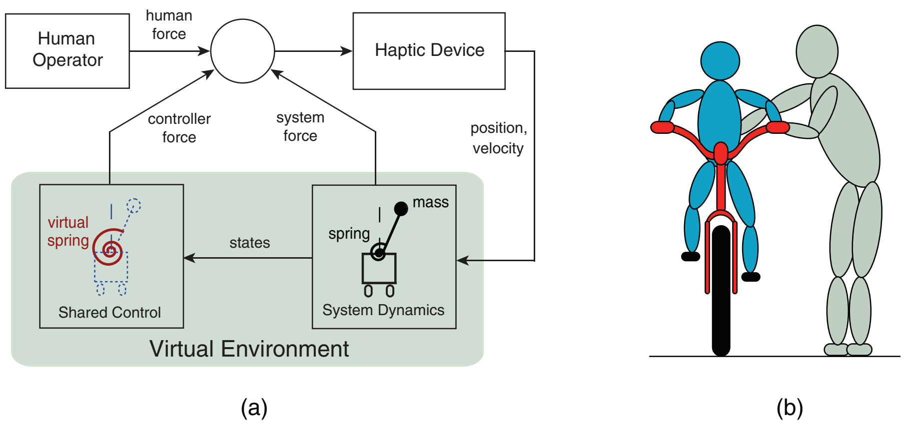

Fig. 1. Figure depicts schematic representation of example shared control implementations for training of bicycle riding task in virtual (a) and physical (b) environments.

and the shared controller is depicted as a spring that helps stabilize this balancing task. Thanks to the stabilizing effects of the shared controller (spring implemented in software), the coupled system (the physical pendulum together with the shared controller) is more stable around the desired fixed point than the physical pendulum is without the shared controller. Figure 1(b) depicts the corresponding scenario in the physical world, where the role of the shared controller is played by the person holding the bicycle to help stabilize the rider.

Shared control with error reduction has consistently been shown to improve task performance in both physical and virtual environments [Griffiths and Gillespie 2005; Yoneda et al. 1999; O'Malley et al. 2006]. There is anecdotal evidence of improvements in skill acquisition with error reducing shared control implemented in a virtual environment training scheme, but clear evidence of the efficacy of error-reduction for virtual training enhancement has not been reported. In one study, Nudehi et al. [2005] implemented shared controllers for training in minimally invasive surgery, but no human subject studies were reported. The authors' implementation of error-reducing shared controller has been shown to affect motor skill acquisition through improved retention trends from one training session to the next compared to practice without assistance [O'Malley et al. 2006], but no strong conclusions regarding training efficacy can be drawn from these data.

This article shows the negative efficacy of fixed-gain error-reducing shared control on training of a manual control task in a haptic virtual environment. In the authors' prior work, error-reducing shared control was shown to be effective at performance enhancement in virtual environments [O'Malley and Gupta 2003]. The extension of prior work to investigate the application of the error-reducing shared controller for training was motivated by evidence from the literature that humans tend to maintain a goal in terms of the kinematics of the end-effector for motor planning and control [Wolpert and Jordan 1995; Shadmehr and Mussa-Ivaldi 1994]. The error-reducing shared controller acts on the user by displaying a force that elicits desired motions of a controlled end-effector. Intuitively, if users plan their motor actions based on desired end-effector behavior, then the authors' error-reducing shared controller should influence skill acquisition and aid the user in coordinating control of the system.

Specifically, a manual skill acquisition experiment is conducted using two different training protocols based on shared control with error reduction, in addition to a control protocol. The protocols include training with haptic guidance (shared control) throughout all assistance subsessions, training with shared control for the first quarter of each assistance subsessions (strategy), and haptic enabled virtual

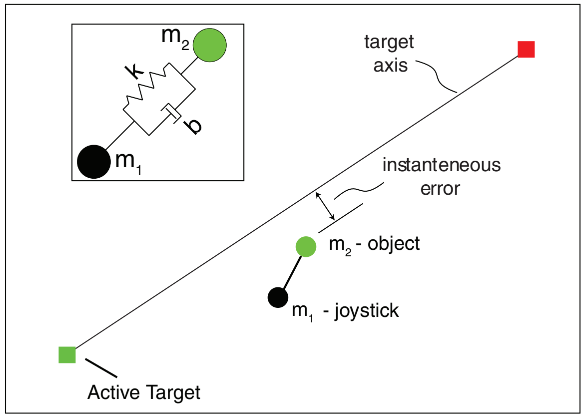

Fig. 2. Target-hitting task: Subjects control location of m_1 (joystick) in order to cause m_2 (object) to hit the desired target. Inset shows virtual underactuated system. The user controls the system by applying forces to mass m_1 through a force feedback joystick based interface. Instantaneous error is defined as the deviation of m_2 (object) from the target axis.

practice (no assistance). The strategy protocol is motivated by the work of Sterr et al. [2002], which demonstrates the importance of training dose in rehabilitation. The authors introduce the strategy treatment since preliminary results from a related study by Li et al. [2006] indicate that different doses of shared control assistance yield significant differences in training performance. In this article, a more thorough study with larger numbers of subjects is conducted to determine the efficacy of shared control for training of a manual control task in a haptic virtual environment. Despite the known benefits of shared control for performance enhancement, the error-reducing shared controller has negative efficacy for training of the manual control task. In addition, while the authors took care to design the haptic guidance scheme based on knowledge of human motor control and the task, results suggest that motor skill acquisition is much more complex that can be intuitively imagined.

The article 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 described in Section 3. Section 4 provides details of the experiment design. The negative efficacy of the fixed-gain error-reducing shared control is presented with experimental results and supporting statistical analysis in Section 5. Section 6 discusses the experimental findings. Finally, Section 7 concludes the article.

2. SYSTEM AND TASK DESCRIPTION

A dynamic manual control task is used to investigate the efficacy of training protocols based on shared control with error reduction. The task consists of controlling a two-mass system, depicted in Figure 2, to alternatively hit targets in the workspace. The dynamics of the controlled system are modeled as two point masses connected by a spring k and damper b in parallel. The system has four degrees of freedom, namely the x and y motion of both masses m_1 and m_2. However, subjects can only directly control 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, the two-mass system is an underactuated system that 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 forces generated by the interactions of the two masses connected by the spring-damper are significant such that subjects can accurately control motion of the system. Haptic feedback has been shown to be an important factor for

Fig. 3. Subject seated at the force feedback joystick, viewing the target-hitting task.

enhancing performance and learning of similar dynamic control tasks [Huang and Gillespie 2007]. In this article, the authors examine the effect of assistance forces from a haptic guidance shared controller during training, which are displayed to the user in addition to the forces of interaction due to the system's inherent dynamics.

2.1 Hardware

An Impulse Engine 2000 joystick from Immersion Inc., shown in Figure 3, is used to provide high-fidelity haptic simulations of the two-mass system. The force feedback joystick has two degrees-of-freedom and a workspace of $6'' \times 6''$. The device exhibits low backdrive friction ($\leq$0.14 N) and high-sensor resolution ($\sim$0.02mm). All simulations run 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 8.9 N in the workspace. The virtual environment graphics are created using OpenGL.

The virtual environment task is rendered using an impedance control mode, where user motion is measured via optical encoders on the joystick, and forces are computed and commanded according to the equations of motion of the system and shared controller. The inherent dynamics of the haptic device are neglected, as is commonly done for high-fidelity impedance type haptic displays. The haptic interface used in the experiments has low friction, is free of backlash, and is highly backdriveable. The device has relatively low inertia and sufficiently high structural stiffness. These features, combined with the relatively low human motion input velocities and absence of impacts in the virtual environment task, enable neglect of the parasitic inertial and Coriolis effects of the device.

2.2 Task

A target-hitting task is used to study manual control of a two-mass underactuated system. Subjects view the virtual environment on a computer monitor and are asked to control the motion of mass m_1 via a two degree-of-freedom haptic device. Through the two-mass system's dynamics, the subjects are able to indirectly control mass m_2 to alternately hit a fixed pair of targets. Figure 3 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 double-mass system. At any given time, one target is active, indicated by changing its color to green. The other is the inactive target, displayed in red. After m_2 contacts the active (green) target, the target colors change to indicate that the previous inactive target (red) is now active (green). Figure 2 illustrates the target pair that is utilized in the experiments. The targets are equidistant from the origin; therefore, the subjects need to move the joystick, directly coupled to m_1, rhythmically, along the sloped path (referred to as the target axis), to cause m_2 to alternately hit the target pair.

3. SHARED CONTROL BASED HAPTIC ASSISTANCE

The goal of the experiments is to investigate efficacy of training protocols based on shared control with error reduction for training of the task described in Section 2. Haptic assistance is provided by additional forces displayed to the subjects via a force feedback joystick. The shared control paradigm for haptic assistance represents active intervention and is implemented through a model based controller. This training scheme is inspired from virtual fixtures but differs from them in that perceptual constraints are not implemented on user input but on user output, and are reflected to the user through the inverse dynamics of the system to be controlled. As stated in the Introduction, the development of this controller was motivated by evidence from the literature that humans tend to maintain a goal in terms of the kinematics of the end-effector for motor planning and control [Wolpert and Jordan 1995; Shadmehr and Mussa-Ivaldi 1994]. In the target-hitting task, m_2 denotes the end-effector. Therefore, following these motor control findings, the control algorithm is designed such that the control effort is applied to m_2 to reduce deviation (error) from the target axis. The forces corresponding to the control actions are fed back to the user through the inverse dynamics of the system. The authors hypothesize that if users indeed plan their motor actions based on the end-effector or dynamics of the external system, then such an implementation of haptic assistance should influence skill acquisition and aid the user in coordinating control of the system.

The dynamics of the task are modeled in order to apply active assistance from the error-reducing shared control algorithm. Defining the x-axis to be aligned with the line connecting the target pair (target axis), and the y axis to be perpendicular to x-axis, the dynamics of the spring-mass system can be described by the following equations of motion:

$$m_1\ddot{x}_1 - F_{kx} = F_x \tag{1}$$
$$m_1\ddot{y}_1 - F_{ky} = F_y \tag{2}$$
$$m_2\ddot{x}_2 + F_{kx} = 0 \tag{3}$$
$$m_2\ddot{y}_2 + F_{ky} = 0 \tag{4}$$

where x_1, y_1, x_2, and y_2 represent the x and y positions of masses m_1 and m_2, while F_{kx} and F_{ky} denote the x and y components of the forces arising from the spring damper pair. The internal forces are calculated as

$$F_{kx} = k(x_2 - x_1) + b(\dot{x}_2 - \dot{x}_1) \tag{5}$$
$$F_{ky} = k(y_2 - y_1) + b(\dot{y}_2 - \dot{y}_1). \tag{6}$$

Finally $F_x = F_{hx} + F_{sx}$ and $F_y = F_{hy} + F_{sy}$ denote the external forces exerted on mass m_1. External forces are due to two sources: F_h components are the portion of the forces exerted by the human operator, whereas F_s components denote the forces exerted through the actuators of the force feedback device due to the shared control assist. Note that in these equations m_1 represents the mass of physical haptic device, while all other parameters (m_2, k and b) are associated with the virtual task dynamics.

For the task described in this article, 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. During shared control assistance, the motion of m_2 is constrained along the target axis. The constraint on the motion of m_2 is derived such that the swing of m_2 normal to the target axis is suppressed. A simple feedback controller is implemented for position control of mass m_2 without explicitly deriving F_{kx} and F_{ky} in system dynamic equations. The desired controlled dynamics for m_2 along the y-axis are defined as

$$\ddot{y}_2 + 2\lambda\dot{y}_2 + \lambda^2 y_2 = -K_p y_2 - K_v \dot{y}_2 \tag{7}$$

where K_v and K_p are control gains rendering the closed loop system dynamics stable. Note that the dynamics along the x-axis are kept unchanged. Effectively, the action of the shared controller is to feed the constraint forces imposed on m_2 to m_1 via the inverse dynamics of the dual mass-spring-damper system described by Equations (1) through (6).

The forces to be displayed due to the shared controller, F_{sx} and F_{sy} can be derived by eliminating $\ddot{x}_2$ and $\ddot{y}_2$ from Equation (7) using Equations (1) through (6), to get

$$F_{sx} = 0 \tag{8}$$

$$F_{sy} = m_1\ddot{y}_1 - m_2[(K_v + 2\lambda)\dot{y}_2 + (K_p + \lambda^2)x_2] \tag{9}$$

Equations (8) and (9) represent the shared control architecture originally derived in O'Malley et al. [2006] and implemented for the target hitting task in this article. This implementation of shared control reduces the difficulty of the task by altering the dynamics of the controlled system to help suppress the motion of m_2 normal to the target axis. Specifically, the shared controller applies forces to decrease perpendicular deviations from the preferred trajectory, forcing the motion of m_2 to stay along the target axis. As discussed in the Introduction, when the shared controller is active, the dynamics of the virtual task are modified along the y-axis, in a way similar to a parent holding an child's bike to help with its stabilization.

Proper implementation of the error reducing shared controller requires full knowledge of the model of the task and haptic interface. The parameters pertaining to the virtual task are user-defined to elicit desired behaviors of the two-mass system and controller. The dynamics attributed to the haptic device are neglected in this study due to the selection of a high-fidelity haptic interface. Such a device exhibits negligible friction and very low effective inertia, and the velocities and accelerations of the haptic device experiences are relatively low. Therefore, the dynamics of the device (specifically its mass m_1) are neglected during the implementation of the controller. Parasitic forces due to the existence of the haptic device and modeling errors exist, but these forces are negligible when compared with the forces rendered through the task dynamics.

4. EXPERIMENT DETAILS

An experiment was conducted to investigate the efficacy of two different training protocols, both based on shared control with error reduction, for manual skill training in virtual environments. Subjects were assigned to three groups (two shared control groups and a control group). The experiment was composed of 11 sessions, including an evaluation session, 9 training sessions, and a retention session. Each session contained three subsessions: preassistance baseline, assistance, and postassistance baseline. Each subsession consisted of 14 trials, with each trial lasting 20 seconds. Details of the experiment design are schematically represented in Figure 4. The control group, also referred to as the no assistance (N) group, received no (force feedback) assistance during assistance subsessions of the experiment. This group represents the case of virtual practice, with haptic and visual feedback of the task and environment. The shared control (A) group was provided with assistance via an error-reducing shared control implementation for the duration of assistance subsessions. The third group, called the strategy (S)

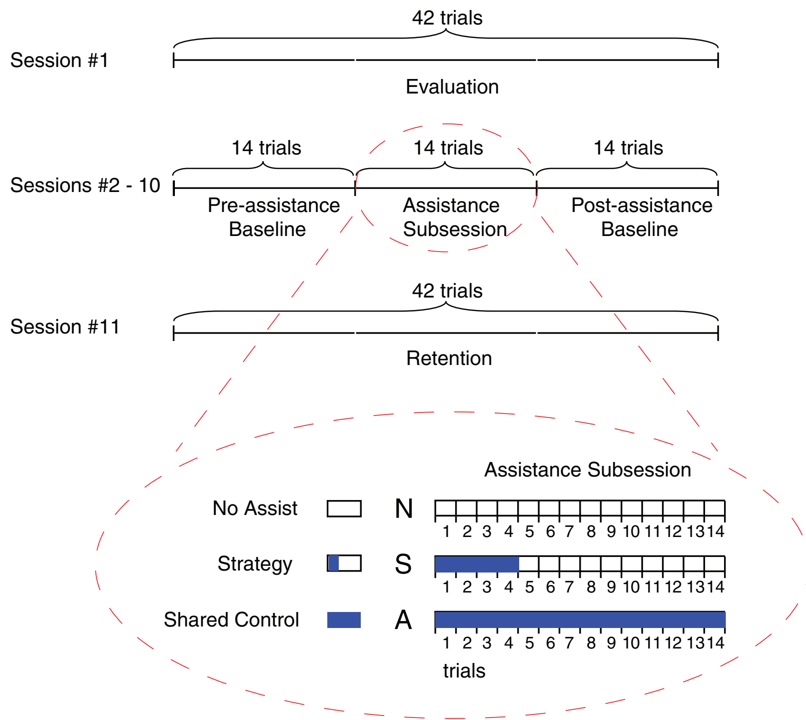

Fig. 4. Figure presents a schematic representation of the experiment design, which consists of one evaluation, nine training and one retention sessions. Each training session contains three subsessions: preassistance baseline, assistance, and postassistance baseline. During each training session, the 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 assistance subsession and the shared control (A) group is provided with shared control assistance throughout all 14 trials of the assistance subsession.

group, received shared control assistance during the first four trials of each assistance subsession. Preassistance and postassistance baseline sessions were used to compare performance of the task across groups without the influence of haptic guidance.

During assistance subsessions, the shared control (A) group was expected to outperform the no assistance (N) group, due to the simplified dynamics introduced by the haptic guidance. Additionally, the authors hypothesized that if the shared control (A) protocol increased training effectiveness, then the performance of the shared control (A) group would outperform the no assistance (N) group in postassistance baseline subsessions. The authors also hypothesized that the strategy (S) protocol would be effective for training, since the error-reducing shared controller would demonstrate the need to limit error. The performance of the strategy (S) group was expected to be better than the no assistance (N) group. Furthermore, the strategy (S) group was expected to outperform the shared control (A) group, since the strategy group has both the advantages of simplified dynamics (strategy) demonstration compared to the no assistance (N) group and greater exposure to the unassisted task dynamics than the shared control (A) group.

Three sets of parameter values for the underactuated system are utilized to increase the difficulty of the task. Table I lists the three selected sets of system parameters that govern the dynamic response of this system. These parameter sets were varied in a controlled manner during the experiment to

Table I. 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

increase the complexity of the task, yet enabling data analysis and comparisons between groups, subjects, and experiment sessions. Within each 14 trial subsession, 4 repetitions of parameter sets 1 and 2, and 3 repetitions of parameter set 3 are presented. The order of presentation is controlled in such a way that the first three trials of every subsession contain one presentation of each set of system parameters. Similarly, the last three trials of every subsession contain one presentation of each set of system parameters. Fixed control parameters used to implement the shared control assistance are selected as $\lambda = 1\, rad/s$, $K_p = 70\, N/m$, and $K_v = 1\, Ns/m$.

Twenty-four subjects (7 female, 17 male, ages 18–25, 2 left-handed), primarily undergraduate students in engineering, participated in the experiments. The handedness of the subjects was not included as a factor in the experiment, since statistical analysis conducted on preliminary experimental data showed that left-handed subjects are not significantly different from the right-handed subjects for the task in consideration. Subjects were instructed to control the motion of mass m_1 via the force feedback joystick, to cause mass m_2 to alternately hit a fixed pair of targets. At the end of each trial, the number of target hits was reported to the subjects.

The 24 subjects were assigned to one of three training protocols based on their initial performance of the target-hitting task. Before the experiment, each subject was given a maximum of 5 minutes to become familiar with the haptic joystick and the virtual task. In order to control individual differences in performance across subjects, each subject was asked to perform the task through an evaluation session, administered without haptic assistance. The purpose of the evaluation session was to measure initial task performance of each subject so that well-balanced group assignments could be made. After the evaluation session, each subject was scored based on the total number of target hits. To account for the change in system parameters, a normalized hit count measure was introduced to account for variations in the system natural frequencies. Subjects were ranked according to their normalized hit count score, and were divided into eight sets with respect to their ranking. Then, subjects from each set were randomly assigned into one of the three paradigms (no assistance [N], strategy [S], and shared control [A]) such that the average score for the three groups was roughly equivalent at the start of training sessions. An analysis of variance (ANOVA) confirmed that the differences among groups was not statistically significant in terms of both normalized hit count and average error metrics after this assignment.

All groups completed nine training sessions. The no assistance (N) group served as the control set with no haptic assistance provided during the assistance subsession. The "no assistance" paradigm is akin to practice. In this training paradigm, subjects felt the forces generated solely due to the internal dynamics of the system. In contrast, during assistance subsessions with the shared controller, subjects felt forces due to both the internal dynamics of the system and the augmented forces intended to assist task completion. The strategy (S) group was provided with shared control assistance in the first 4 of 14 trials of the assistance subsessions, whereas the shared control (A) group was provided with shared control assistance throughout all 14 trials of the assistance subsessions. In order to assess the improvement of subjects across the nine training sessions, baseline subsessions of 14 trials administered without assistance were completed before and after each assistance subsession. An assistance subsession and its corresponding pre- and postassistance baseline subsessions took place in a single sitting. The nine

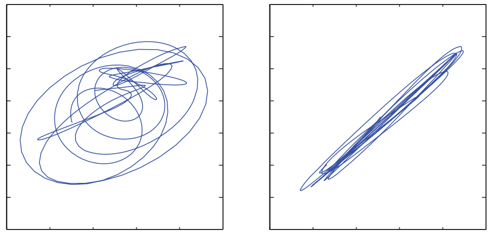

Fig. 5. Figures depict trajectories of mass m_2 during the evaluation and the retention sessions for a typical subject. Subfigure on the left (a) presents typical trajectories during evaluation sessions, whereas the subfigure on the right (b) presents typical trajectories during the retention sessions.

training sessions were separated by 2 to 3 days, such that subjects completed all sessions in no less than 3 but no more than 4 weeks. One month after the final training session, all subjects were recalled to complete one retention session. In the retention session, no haptic assistance was provided as was done for the evaluation session.

To summarize, the experiment consisted of three factors, namely assistance mode, session, and parameter set. Assistance mode was between-subjects, with levels no assistance (N), strategy (S), and shared control (A). Session was within-subjects, with levels evaluation, training (9 in all), and retention, for a total of 11 levels. Parameter set was also within-subjects, with three possible sets.

5. RESULTS

Two performance measures are used to assess subject performance for the target-hitting task, namely normalized hit count and the average error. Normalized hit count is the total number of target hits within one 20-second trial normalized by the natural frequency of the corresponding dynamic system. The average error is the average of the instantaneous position deviation of the mass m_2 from the target axis. A schematic representation of instantaneous position deviation is depicted in Figure 2. Together, these performance measures capture the features of the task, where normalized hit count gives an assessment of speed of execution, while average error monitors the ability of the subject to maintain a trajectory along the target axis. Average error is treated as a secondary performance metric, since subjects are not instructed to reduce the deviation of m_2 from the target axis. However, this metric is of interest because it helps investigate if shared control assistance, designed to reduce the deviation from target axis, conveys this goal as a strategy to subjects. Moreover, average error can be used to compare the training effectiveness of the shared control protocol to the strategy protocol.

Figure 5 presents trajectories of mass m_2 during the evaluation (a) and the retention (b) sessions for a typical subject. All subjects adopted an excitation strategy such that their trajectories converged to a straight-line path between targets. While not explicitly stated as a goal of task performance, such trajectories in general enable higher normalized hit count scores.

Figures 6 through 8 show results for normalized hit count and average error for preassistance baseline, postassistance baseline, and assistance subsessions, respectively. As mentioned earlier, haptic assistance due to the shared control algorithm is active only during assistance subsessions for those subjects assigned to groups receiving haptic guidance. For pre- and postassistance subsessions, subjects perform the task without the addition of any control action, allowing for performance comparisons across groups. In the figure, results are grouped based on the training protocol to which subjects were assigned. Normalized hit count follows an increasing trend for all groups, while average error data is identified by its decreasing trend throughout training. Error bars indicate standard errors for the

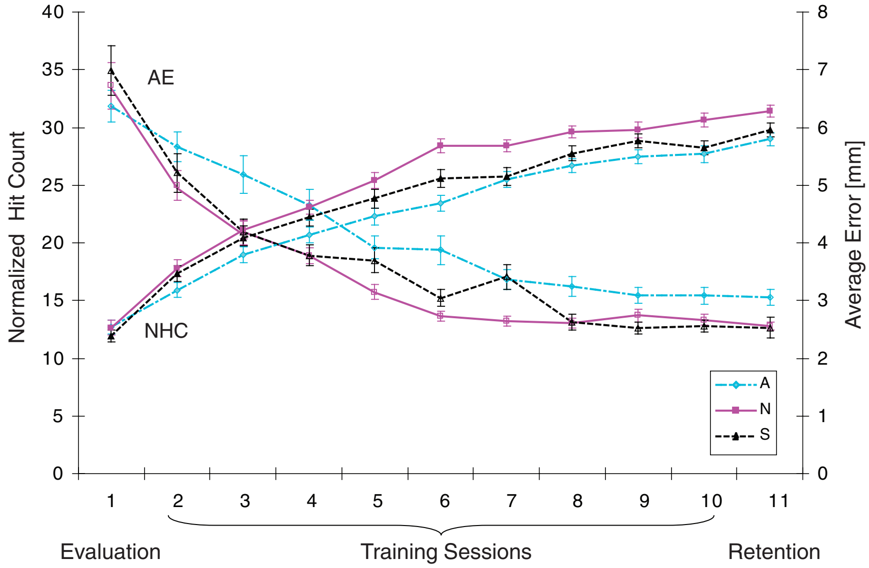

Fig. 6. Preassistance baseline normalized hit count and average error plots for different groups over 11 sessions. Group A represents the shared control group, N is the no assistance group, and S is the strategy group.

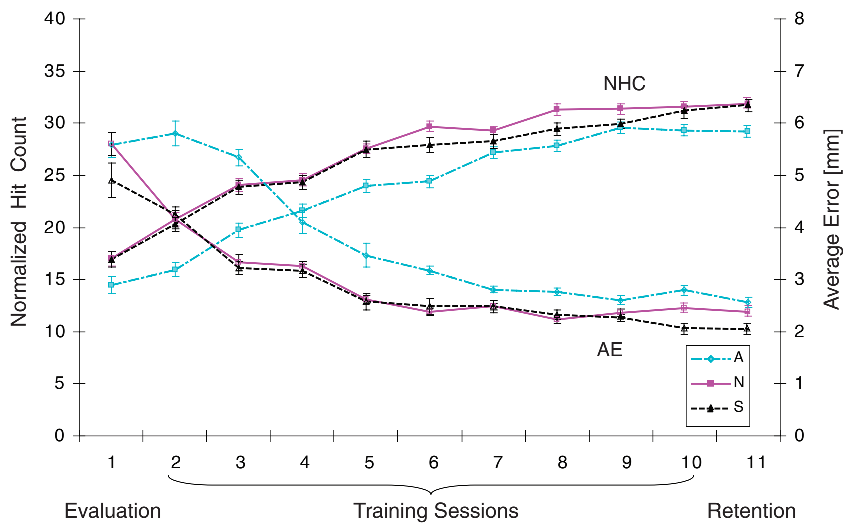

Fig. 7. Postassistance baseline normalized hit count and average error plots for different groups over 11 sessions.

results. In the following figures, N represents the no assistance group, A represents the shared control group, and S represents the strategy group.

Figure 6 presents preassistance baseline normalized hit count and average error results for each training group for the 11 experiment sessions, including the evaluation session (session 1), 9 training sessions (sessions 2–10), and the retention session (session 11). Results show that for the evaluation session, all three groups start at approximately the same performance level in terms of both normalized

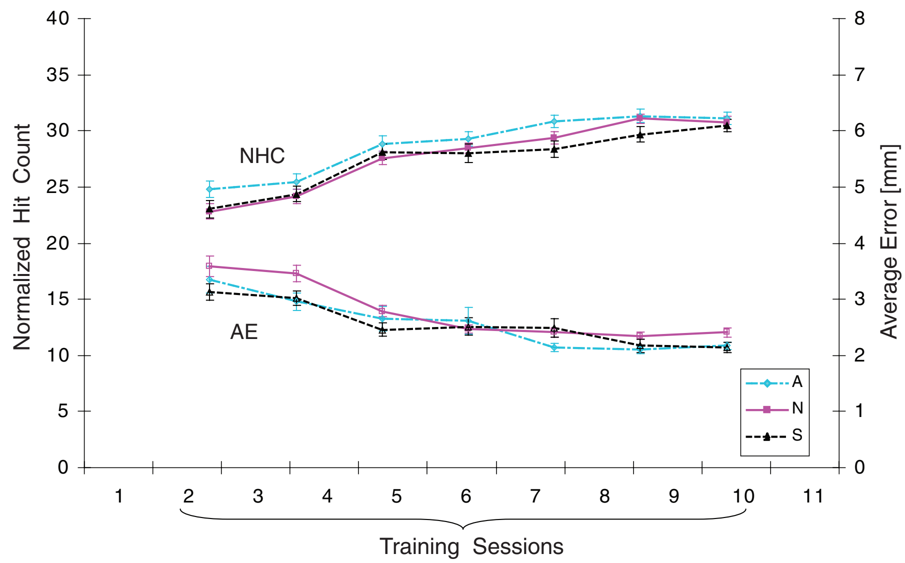

Fig. 8. Assistance subsession normalized hit count and average error plots for different groups over nine training sessions (sessions 2–10).

hit count and average error. This result is validated through a simple one-way ANOVA (group), and the differences are not significant in terms of normalized hit count ([$F(2, 1005) = 2.06$, $p = 0.1284$]) or in terms of average error ([$F(2, 1005) = 0.26$, $p = 0.7747$]). Figure 7 presents postassistance baseline normalized hit count and average error for each training group over 11 sessions.

During the nine training sessions, subjects for all training protocols demonstrate improving performance, as illustrated by increasing values in normalized hit count and decreasing values in average error metrics, for both pre- and postassistance baseline measurements (see Figures 6 and 7). The retention data also show that for all groups, learning continues even 1 month after the last training session (session 10). It is worth noting that average error, the secondary performance measure, decreases even for the no assistance (N) group, which is never exposed to the preferred strategy through the action of the shared controller.

The learning effect over the course of the experiment is significant, starting from approximately 13 normalized hits per trial and improving to 30 normalized hits in the retention session. The improvements in performance are statistically significant between evaluation and retention sessions ([$F(1{,}670) = 1029.4$, $p < 0.0001$] for A, [$F(1{,}670) = 1{,}252.42$, $p < 0.0001$] for N, [$F(1{,}670) = 1{,}214.95$, $p < 0.0001$] for S) calculated using a simple one-way ANOVA (session) in terms of normalized hit count for preassistance baseline and postassistance baseline. The performance of different groups measured by average error parallels that of normalized hit count, exhibiting a decrease in average error as learning takes place. The learning effect at the end of the experiment is significant for all three groups in term of the average error metric, with [$F(1{,}670) = 415.27$, $p < 0.0001$] for A, [$F(1{,}670) = 266.21$, $p < 0.0001$] for N, [$F(1{,}670) = 286.44$, $p < 0.0001$] for S calculated using a simple one-way ANOVA (session) in terms of average error for preassistance baseline and postassistance baseline.

As depicted in Figures 6 through 7, performance of all groups in terms of both performance metrics saturates at a similar level. There is no statistically significant difference in the performance of training protocol groups for the last two sessions of the experiment for either performance measure. However, the overall performance of the groups is significantly different for preassistance, postassistance baseline,

Table II. Summary of Significance Measured by ANOVA in Terms of Normalized Hit Count

Effect	Preassistance	Assistance	Postassistance
Group	$F(2{,}2997) = 40.06,\ p < 0.0001^*$	$F(2{,}2997) = 11.73,\ p < 0.0001^*$	$F(2{,}2997) = 72.85,\ p < 0.0001^*$
Session	$F(8{,}2997) = 111.7,\ p < 0.0001^*$	$F(8{,}2997) = 112.83,\ p < 0.0001^*$	$F(8{,}2997) = 119.6,\ p < 0.0001^*$
Interaction	$F(16{,}2997) = 0.71,\ p = 0.79$	$F(16{,}2997) = 0.6,\ p = 0.8847$	$F(16{,}2997) = 1.69,\ p = 0.0408^*$

Table III. Summary of Significance Measured by ANOVA in Terms of Average Error

Effect	Preassistance	Assistance	Postassistance
Group	$F(2{,}2997) = 36.61,\ p < 0.0001^*$	$F(2{,}2997) = 7.95,\ p = 0.0004^*$	$F(2{,}2997) = 88.26,\ p < 0.0001^*$
Session	$F(8{,}2997) = 64.85,\ p < 0.0001^*$	$F(8{,}2997) = 55.01,\ p < 0.0001^*$	$F(8{,}2997) = 78.93,\ p < 0.0001^*$
Interaction	$F(16{,}2997) = 1.27,\ p = 0.2088$	$F(16{,}2997) = 1.14,\ p = 0.3065$	$F(16{,}2997) = 4.91,\ p < 0.0001^*$

Table IV. Summary of Significance Measured by Differences of Least Square Means in Terms of Normalized HitCount

Group Comparison	Preassistance	Assistance	Postassistance
A vs. N	$p < 0.0001^*$	$p > 0.05$	$p < 0.0001^*$
N vs. S	$p = 0.0024^*$	$p > 0.05$	$p = 0.0566^*$
A vs. S	$p = 0.0289^*$	$p = 0.0388^*$	$p < 0.0001^*$

Table V. Summary of Significance Measured by Differences of Least Square Means in Terms of Average Error

Group Comparison	Preassistance	Assistance	Postassistance
A vs. N	$p < 0.0001^*$	$p = 0.0325^*$	$p < 0.0001^*$
N vs. S	$p > 0.05$	$p = 0.0222^*$	$p > 0.05$
A vs. S	$p = 0.0003^*$	$p > 0.05$	$p < 0.0001^*$

and assistance subsession throughout all training sessions (session 2 to session 10). A mixed design two-way (group, session) ANOVA was carried out to determine significance of results for the three groups. The results revealed a significant main effect of group and session for preassistance, postassistance, and assistance subsessions in terms of both normalized hit count and average error. A summary of these two-way ANOVA results is listed in Table II and Table III.

It is worth noting that the interaction effect is significant for postassistance baseline for normalized hit count ([$F(16{,}2997) = 1.69$, $p = 0.0408$] and average error ([$F(16{,}2997) = 4.91$, $p < 0.0001$]) but not significant for preassistance baseline and assistance subsessions. The significant interaction is probably due to the influence of the assistance subsessions training protocols, which will be discussed in detail in Section 6.

In order to further explore the influence of training protocols with shared control on the performance of each group, the difference of least square means (LSM) is used. This statistical analysis method uses an adjusted mean for each group that isolates the effect of each individual group, then provides specific comparisons between each combination of two groups. A summary of all pertinent comparisons for the least square means analysis is listed in Tables IV and V. The investigation of least square means for assistance subsessions reveals that the strategy (S) group is significantly worse than the shared control (A) group in terms of normalized hit count ($p = 0.0388$). The no assistance (N) group is significantly worse than both shared control (A) and strategy (S) groups in terms of average error ($p = 0.0325$ for A vs. N, and $p = 0.0222$ for S vs. N). These results support the hypothesis that shared control enhances performance of the target-hitting task by reducing the average error.

While LSM analysis reveals the positive influence of shared control on performance during assistance subsessions, the haptic guidance protocols do not demonstrate efficacy when subjects perform the task without assistance, as in the pre- and postassistance baseline subsessions. The difference of least square means statistical analysis indicates significant differences for the following combinations between different training protocols in baseline subsessions (shown in Tables IV and V: the No assistance (N) group performs significantly better than the shared control (A) group in terms of both normalized hit count and average error with $p < 0.0001$ for preassistance baseline and with $p < 0.0001$ for postassistance baseline over nine training sessions. Statistical analysis of the no assistance (N) group compared to the strategy (S) group shows a similar trend with the (N) group performing significantly better than the strategy (S) group in terms of normalized hit count for preassistance baseline ($p = 0.0024$) and postassistance baseline ($p = 0.0566$) subsessions. When considering average error, there is no significant difference in measures between the no assistance (N) and strategy (S) groups. When comparing the two training groups that experienced shared control, the strategy (S) group performs significantly better than the shared control (A) group in terms of both normalized hit count ($p = 0.0289$) and average error ($p = 0.0003$) for preassistance baseline subsessions. For postassistance baseline subsessions, the strategy (S) group outperforms the shared control (A) group in both performance metrics ($p < 0.0001$).

6. DISCUSSION

During assistance subsessions, performance enhancement is expected for those subjects experiencing shared control with error reduction. Indeed, results from this experiment demonstrate that the performance of the shared control (A) group during assistance subsessions is significantly improved over performance during pre- and postassistance baselines, when the shared control is inactive. According to a repeated measures one-way (session) ANOVA, normalized hit count and average error measures are significantly different ($p < 0.001$). This result is in good agreement with previous studies in the literature that demonstrate the effectiveness of haptic guidance for performance enhancement [Griffiths and Gillespie 2005; Yoneda et al. 1999; O'Malley et al. 2006].

The primary goal of this study was to determine the efficacy of haptic guidance, implemented via an error-reducing shared control algorithm, on training of a manual target-hitting task. Despite performance gains during assistance subsessions, the gains are not maintained when haptic guidance is removed. There is in fact a negative training efficacy of fixed-gain error-reducing shared control, noted by the lagging performance of the shared control (A) group compared to that of the virtual practice (N) group for baseline subsessions. In Figures 6 and 7, results show that the shared control (A) group performs significantly worse than the no assistance (N) group during both pre- and postassistance baseline subsessions ($p < 0.0001$) in terms of normalized hit count. The shared control (A) protocol as implemented, with assistance provided throughout all assistance subsession trials, is not effective for training.

To further explore the effects of the shared control (A) protocol on hit count performance, learning trends within a session and across sessions are analyzed in Figure 9. Figure 9 consists of two plots. The first is a line plot displaying absolute task performance improvement within a session, with each line segment corresponding to one training session. The second is a bar plot representing the percent change of performance within a session. For both plots, the data represents a comparison of the average performance of the last three trials of the preassistance baseline subsession to the average performance of the first three trials of postassistance baseline subsessions. The average values presented are of interest because they characterize within session performance of a group just before and just after the assistance subsession, quantifying the amount of learning that occurred within the assistance subsession. By averaging three trials, task performance for each parameter set is captured due to the constraints implemented in experiment design.

As depicted in Figure 9, the shared control (A) and no assistance (N) groups start from approximately the same performance level in terms of normalized hit count at session 2, just after the evaluation session. However, the slope of the within session learning curve becomes negative for the shared control (A) group after the first assistance subsession. Group (A) experiences fixed-gain error-reducing shared

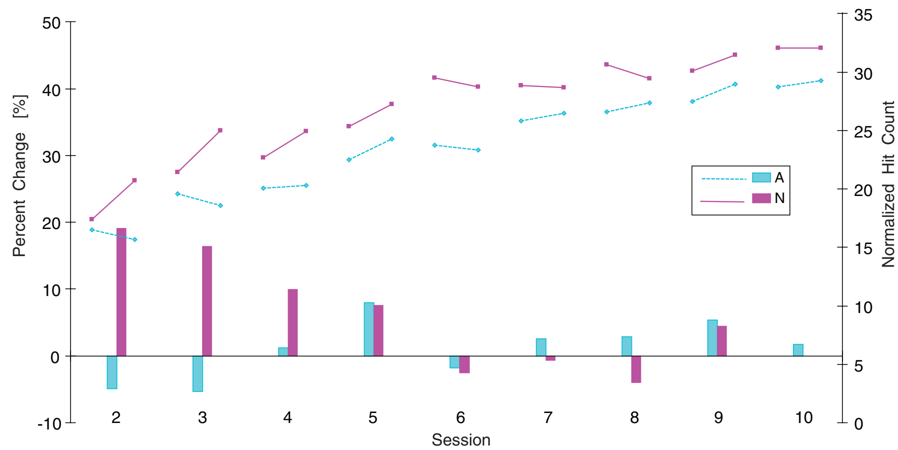

Fig. 9. Absolute task performance (line) and performance change percentage (bar) plots in terms of normalized hit count for shared control (A) and no assistance (N) groups over nine training sessions (sessions 2–10).

control during assistance subsession 2. The effect of exposure to this form of haptic guidance resulted in a *negative* 5% change in performance between the end of the preassistance baseline subsession and the beginning of the postassistance baseline subsession. Comparatively, learning takes place for the no assistance (N) group during the same assistance subsession, resulting in a positive 19% change in performance between the end of the preassistance baseline subsession and the beginning of the postassistance baseline subsession. The negative learning trend continued for the shared control (A) group in session 3, while the no assistance (N) group further enhanced their performance with positive gains. Starting from session 4, the change in performance for the shared control (A) group during assistance subsessions becomes positive.

The existence of negative learning in early sessions of training (sessions 2 and 3) for the shared control (A) group indicates that shared control assistance for the duration of assistance subsessions significantly interferes with learning. However, subjects quickly (starting from session 4) adapt to the display of coupled dynamics introduced by the shared controller. After adaptation, subjects exhibit positive learning trends for the remainder of the training protocol. The negative learning effects within the early sessions as shown in Figure 9 may be attributed to interference in motor learning. Interference and consolidation are observed phenomena that occur when one interacts with secondary dynamic tasks while trying to learn a primary dynamic task [Shadmehr et al. 1995; Brashers-Krug et al. 1996]. In the training paradigms tested here, the shared controller modifies and augments the system dynamics in order to assist in completion of the task. Consequently, the assisted task is a secondary task to be learned. Interference and consolidation are closely affected by the similarity of the stimulus-response mapping between primary dynamics and secondary dynamics, and the greater the differences between primary and secondary tasks, the more severe the interference becomes. It has also been reported that interference decreases with increased learning of the primary dynamics [Siipola and Israel 1993], which supports the observed learning trend of the shared control (A) group. In the literature, increasing the amount of practice time for primary and secondary dynamics or introduction of consolidation time between practice sessions with two distinct dynamics are suggested as solutions to reduce such interference effects [Shadmehr et al. 1995; Brashers-Krug et al. 1996].

Within session, learning trends for the normalized hit count performance measure for the strategy (S) and no assistance (N) groups are illustrated in Figure 10. In these plots, the average of the last three

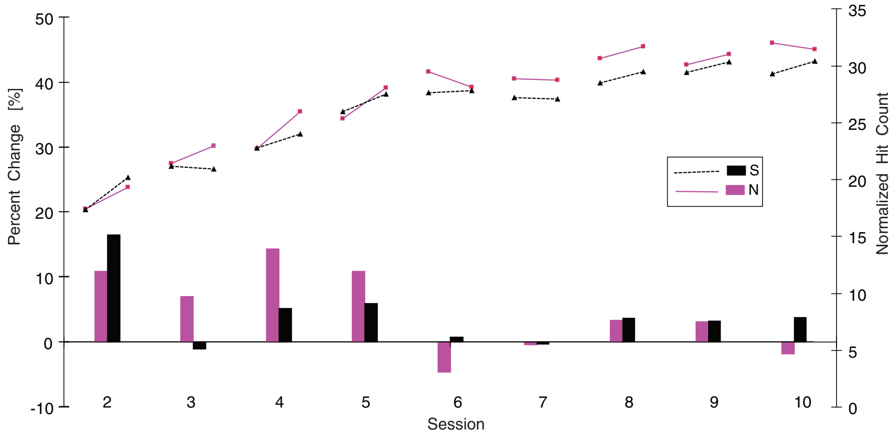

Fig. 10. Absolute task performance (line) and performance change percentage (bar) plots in terms of normalized hit count for strategy (S) and no assistance (N) groups over nine training sessions (sessions 2–10).

trials of the assistance subsession is used in place of the first three trials of the postassistance baseline subsession, as in the comparisons of groups (A) and (N). The last three assistance subsession trials are compared to the last three trials of the preassistance baseline subsession in order to observe the learning effects.

Unlike the shared control (A) protocol, the strategy (S) protocol does not cause subjects to exhibit strong negative learning effects during the early sessions. Hence, the strategy (S) protocol, which provides shared control assistance for only four trials at the beginning of each assistance subsession, does not significantly interfere with learning. One possible explanation may be that the reduced amount of exposure to the shared control algorithm reduced the interference effect. This may have been achieved through extended practice with the unmodified dynamics, as suggested in Siipola and Israel [1993]. As depicted in Figure 10, as well as in Figures 6 through 7, the strategy (S) group follows similar performance trends to that of the no assistance (N) group from session 2 to session 5. However, starting from session 6, the performance of the strategy (S) group consistently lags performance of the no assistance (N) group. The strategy (S) and no assistance (N) protocols are essentially the same except for 4 trials of shared control provided to the strategy (S) group out of 42 total trials per session. Therefore, the negative impact of the strategy (S) protocol on performance, especially during the later sessions of training, may be attributed to the cumulative effect of the strategy (S) group having four fewer trials of practice with the actual task dynamics than the no assistance (N) group at every session. Further experimentation will be required to determine the validity of this hypothesis. The results indicate that the strategy (S) protocol is not as effective as virtual practice for training in terms of the normalized hit count performance measure.

It can be concluded from the analysis of Figures 9 and 10 that both of the shared control based training protocols (A and S groups) exhibit negative efficacy, with hit count performances significantly worse than the virtual practice (N) group. One possible reason for the negative results maybe due to poor design of the shared controller. According to Lintern [1991], during training, a task should be simplified only if the important perceptual invariants of the task are preserved. In the current implementation, the shared controller applies forces to decrease perpendicular deviations from the preferred trajectory, forcing the motion of m_2 to stay along the active target axis. However, a subsequent task analysis by the authors concludes that this target-hitting task is essentially composed of two important aspects: the

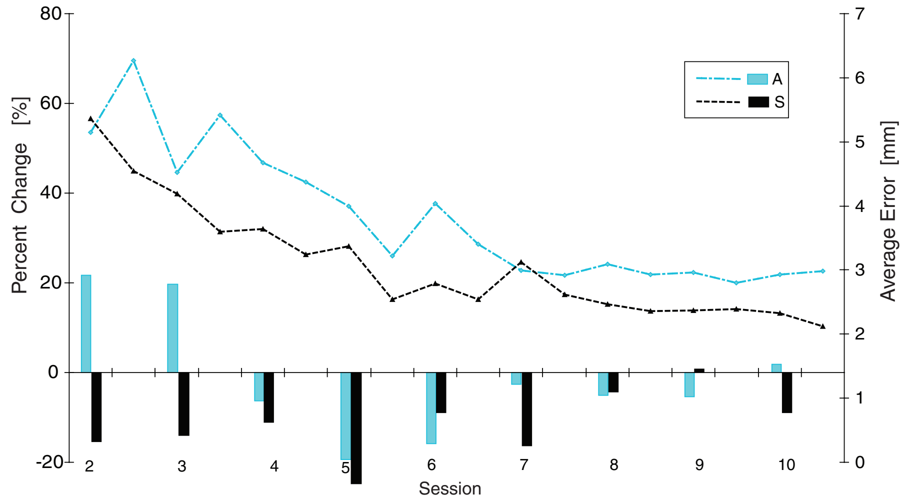

Fig. 11. Absolute task performance (line) and performance change percentage (bar) plots in terms of average error for shared control (A) and strategy (S) groups over nine training sessions (sessions 2–10).

temporal control along the target axis in order to excite the system near its resonant frequency, thus increasing hit count, and the position control perpendicular to target axis, required to minimize error and enable more accurate trajectory control through the targets. Hence, in the current implementation of shared control, only the position control aspects of the task, minimizing error, are exhibited by the shared controller. Therefore, one way to counteract the negative efficacy of the training paradigms may be to redesign the shared controller to capture all critical aspects of the manual control task. In particular, instead of following an intuitive approach to controller design, a careful study of important aspects of the task or analysis of techniques exhibited by high-performing subjects may lead to effective controller designs. For the target-hitting task, a new shared controller design approach is proposed which decomposes the task into two subtasks. The first subtask aims to reduce the deviation from target axis. The second subtask aims to excite the virtual two-mass system near its resonant frequency, enabling higher hit count scores. In order to be effective for training, the authors hypothesize that a shared control algorithm for this task must address both subtasks.

Within session learning trends in terms of average error for the shared control (A) and strategy (S) groups are illustrated in Figure 11. In this figure, the average error measures of the last three trials of the preassistance baseline subsession and the first three trials of the postassistance baseline subsession are used. The shared control (A) group exhibits approximately 20% negative learning at the early stages of training in term of average error, similar to the within session learning curve in term of normalized hit count, while the strategy (S) group exhibits consistent positive learning effects throughout all training sessions. From this observation, the authors conclude that different doses of shared control assistance results in different learning effects, even in terms of the secondary performance measure.

The results show that while exposure to haptic guidance throughout all of training exhibits negative learning effects, the protocol with a shorter duration of exposure (strategy) still shows a performance inferior to that of virtual practice. Hence, the authors conclude that even though the dose effects of shared control during training are significant, no training enhancement can be gained from shared control with error reduction as implemented in this experiment, regardless of the amount of exposure to the haptic assistance. In addition to not capturing all important features of the target-hitting task, the authors attribute the inferior performance of the tested shared control protocols to the way controller is implemented. Fixed controller gains used during assistance may be a major contributor to the negative learning effects of the shared control (A) group and the overall poor performance of both (A) and (S) groups. With fixed control gains for the shared controller, the amount of assistance provided to the subjects stays constant regardless of their performance level. As a consequence, the subjects learn to rely on the existence of the shared controller, rather than developing the required motor skills for successful completion of task in the absence of haptic guidance. Supporting evidence for such phenomena exists in visual augmented feedback-based training of pilots, in which it is reported that learning transfers to the real task *only* if the augmented feedback is presented during periods of large deviation, rather than being presented continually [Lintern et al. 1990]. Similar effects are also reported for rehabilitation in Kahn et al. [2004].

Thus, the authors hypothesize that an adaptive shared control-based training paradigm, which adjusts its control gains based on subject performance, can promote learning. In other words, an adaptive shared control algorithm may increase training effectiveness when compared to virtual practice by exposing subjects to an appropriate amount of haptic assistance based on their performance, such that as performance improves, haptic guidance subsides and eventually is removed entirely. The proposed implementation of an adaptive shared controller follows from the early studies of Bernstein [1967], who proposed that learning can be facilitated by reducing the degrees of freedom of a complex task. The core idea of his hypothesis is inspired from well-developed multiphase 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 also been shown to be the most effective training mechanism for neural networks [Ivanchenko and Jacobs 2003].

7. CONCLUSION

This paper presents experimental results that indicate negative efficacy of haptic guidance for training of a manual control task in a haptic and visual virtual environment. In the study, a target-hitting task was presented to subjects over a month-long period, and performance in the absence of haptic guidance was compared for three training protocols. Two of the protocols incorporated a fixed-gain error-reducing shared controller, which was designed to simplify task dynamics by minimizing motion of the controlled end-effector perpendicular to an axis between targets. The two shared control protocols varied in the duration of exposure to the shared controller. A third protocol enabled virtual practice, where subjects were able to interact with the two-mass system with visual and haptic feedback of the task dynamics but with no haptic guidance. The dose effects of shared control on performance were significant, but results indicate that the regardless of duration of exposure to the shared control algorithm, performance of subjects who were exposed to haptic guidance was inferior to performance of the virtual practice group. Despite knowledge of the task and consideration of human motor control and learning theories when designing the shared control algorithm, results suggest that an intuitive approach to haptic guidance design for training may not be effective and that acquisition of motor skill is a complex phenomenon.

ACKNOWLEDGMENTS

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 and to anonymous reviewers for their insightful comments.

REFERENCES

ADAMS, R. J., KLOWDEN, D., AND HANNAFORD, B. 2001. Virtual traning for a manual assembly task. *Haptics-e 2*, 2.

BERNSTEIN, N. 1967. *The Coordnation and Regulation of Movements*. Oxford/Pergamon.

BETTINI, A., MARAYONG, P., LANG, S., AND OKAMURA, A. M. 2002. Vision assisted control for manipulation using virtual fixtures: Experiments at macro and micro scales. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 3354–3361.

BRASHERS-KRUG, T., SHADMEHR, R., AND BIZZI, E. 1996. Consolidation in human motor memory. *Nature*, 252–255.

EMKEN, J. L. AND REINKENSMEYER, D. J. 2005. Robot-enhanced motor learning: Accelerating internal model formation during locomotion by transient dynamic amplification. *IEEE Trans. Neural Syst. Rehab. Engin. 13*, 1, 33–39.

FEYGIN, D., KEEHNER, M., AND TENDICK, F. 2002. Haptic guidance: Experimental evaluation of a haptic training method for perceptual motor skill. In *Proceedings of the IEEE International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 40–47.

GAMBERINI, L. 2000. Virtual reality as a new research tool for the study of human memory. *CyberPsych. Behav. 3*, 3, 337–342.

GILLESPIE, R. B., S. O'MODHRAIN, P. T., PHAM, C., AND ZARETSKY, D. 1998. Virtual teacher. In *Proceedings of the ASME International Mechanical Engineering Congress and Exposition*. 3354–3361.

GRIFFITHS, P. AND GILLESPIE, R. B. 2005. Sharing control between human and automation using haptic interface: Primary and secondary task performance benefits. *Human Factors 47*, 3, 574–590.

HAANPAA, D. P. AND ROSTON, G. P. 1997. An advanced haptic system for improving man-machine interfaces. *Computers & Graphics 21*, 4, 443–449.

HENMI, K. AND YOSHIKAWA, T. 1998. Virtual lesson and its application to virtual calligraphy system. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 1275–1280.

HUANG, F. AND GILLESPIE, R. B. 2007. Visual and haptic feedback contribute to tuning and online control during object manipulation. *J. Motor Behav. 39*, 179–193.

IVANCHENKO, V. AND JACOBS, R. A. 2003. A developmental approach aids motor learning. *Neural Comput. 15*, 2051–2065.

KAHN, L. E., RYMER, W. Z., AND REINKENSMEYER, D. J. 2004. Adaptive assistance for guided force training in chronic stroke. In *Engineering in Medicine and Biology Society*. Vol. 1. 2722–2725.

KIKUUWE, R. AND YOSHIKAWA, T. 2001. Haptic display device with fingertip presser for motion/force teaching to human. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 868–873.

KOZAK, J. J., HANCOCK, P. A., ARTHUR, E. J., AND CHRYSLER, S. T. 1993. Transfer of training from virtual reality. *Ergonomics 36*, 1, 777–784.

LI, Y., PATOGLU, V., AND O'MALLEY, M. K. 2006. Shared control for training in virtual environments: Learning through demonstration? In *EuroHaptics 2006*.

LINTERN, G. 1991. An informational perspective on skill transfer in human-machine systems. *Human Factors 33*, 3, 251–266.

LINTERN, G. AND ROSCOE, S. N. 1980. *Visual cue Augmentation in Contact Flight Simulation*. Aviation Psychology. Iowa State University Press.

LINTERN, G., ROSCOE, S. N., KOONCE, J. M., AND SEGAL, L. D. 1990. Display principles, control dynamics, and environmental factors in pilot training and transfer. *Human Factors 32*, 3, 299–317.

MASSIMINO, M. J. AND SHERIDAN, T. B. 1994. Teleoperator performance with varying force and visual feedback. *Human Factors 36*, 1, 145–157.

MEECH, J. F. AND SOLOMONIDES, A. E. 1996. User requirements when interacting with virtual objects. *Virtual Reality—User Issues 68*, 1, 1–3.

NUDEHI, S. S., MUKHERJEE, R., AND GHODOUSSI, M. 2005. A shared control approach to haptic interface design for minimally invasive telesurgical training. *IEEE Trans. Control Syst. Tech. 13*, 4, 588–592.

O'MALLEY, M. K. AND GUPTA, A. 2003. Passive and active assistance for human performance of a simulated underactuated dynamic task. In *Proceedings of the IEEE International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. 348–355.

O'MALLEY, M. K., GUPTA, A., GEN, M., AND LI, Y. 2006. Shared control in haptic systems for performance enhancement and training. *ASME J. Dynamic Syst. Measur. Control* (Special Issue on Novel Robotics and Control) *128*, 1, 75–85.

O'MALLEY, M. K., HUGHES, K. J., MAGRUDER, D. F., AND AMBROSE, R. O. 2003. Simulated bilateral teleoperation of robonaut. In *Proceedings of the AIAA Space Conference*.

PATTON, J. L. AND MUSSA-IVALDI, F. A. 2004. Robot-asisted adaptive training: Custom force fields for teaching movement patterns. *IEEE Trans. Biomed. Engin. 51*, 4, 636–646.

Patton, J. L., Mussa-Ivaldi, F. A., and Rymer, W. Z. 2001. Altering movement patterns in healthy and brain-injured subjects via custom designed robotic forces. In *Proceedings of the IEEE International Conference of Engineering Medicine Biology Society*. 1356–1359.

Richard, P. and Coiffet, P. 1995. Human perceptual issues in virtual environments: sensory substitution and information redundancy. In *Proceedings of the IEEE International Workshop in Robot and Human Communication*. 301–306.

Rosenburg, L. B. 1993. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *Proceedings of the IEEE Annual Virtual Reality International Symposium*. 76–82.

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

Shadmehr, R. and Mussa-Ivaldi, F. A. 1994. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci. 14*, 3208–3224.

Siipola, E. M. and Israel, H. E. 1993. Habit interferences as dependent upon stage of training. *Amer. J. Psych. 45*, 205–227.

Sterr, A., Elbert, T., Berthold, I., Klbel, S., Rockstroh, B., and Taub, E. 2002. Longer versus shorter daily constraint-induced movement therapy of chronic hemiparesis: An exploratory study. *Archives Physical Med. Rehab. 83*, 10, 1374–1377.

Todorov, E., Shadmehr, P., and Bizzi, E. 1997. Augmented feedback presented in a virtual environment accelerates learning of a difficult motor task. *J. Motor Behav. 29*, 2, 147–158.

Williams, L. E. P., Loftin, R. B., Aldridge, H. A., Leiss, E. L., and Bluethmann, W. 2002. Kinesthetic and visual force display for telerobotics. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 1249–1254.

Wolpert, D. M. and Jordan, Z. G. M. I. 1995. Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Exper. Brain Resear. 103*, 3, 460–470.

Yokokohji, Y., Hollis, R., Kanade, T., and Henmi, K. 1996. Toward machine mediated training of motor skill. In *Proceedings of the IEEE International Workshop on Robot and Human Communication*. 32–37.

Yoneda, M., Arai, F., Fukuda, T., Miyata, K., and Naito, T. 1999. Assistance system for crane operation with haptic display: Operational assistance to suppress round payload swing. In *Proceedings of the IEEE International Conference on Robotics and Automation*. 2924–2929.

Received August 2006; revised May 2007, November 2007; accepted December 2007