Shared Control in Haptic Systems for Performance Enhancement and Training

Marcia K. O'Malley
e-mail: omalleym@rice.edu

Abhishek Gupta
e-mail: abhi@rice.edu

Matthew Gen
e-mail: mgen@rice.edu

Yanfang Li
e-mail: yvonneli@rice.edu

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

This paper presents a shared-control interaction paradigm for haptic interface systems, with experimental data from two user studies. Shared control, evolved from its initial telerobotics applications, is adapted as a form of haptic assistance in that the haptic device contributes to execution of a dynamic manual target-hitting task via force commands from an automatic controller. Compared to haptic virtual environments, which merely display the physics of the virtual system, or to passive methods of haptic assistance for performance enhancement based on virtual fixtures, the shared-control approach offers a method for actively demonstrating desired motions during virtual environment interactions. The paper presents a thorough review of the literature related to haptic assistance. In addition, two experiments were conducted to independently verify the efficacy of the shared-control approach for performance enhancement and improved training effectiveness of the task. In the first experiment, shared control is found to be as effective as virtual fixtures for performance enhancement, with both methods resulting in significantly better performance in terms of time between target hits for the manual target-hitting task than sessions where subjects feel only the forces arising from the mass-spring-damper system dynamics. Since shared control is more general than virtual fixtures, this approach may be extremely beneficial for performance enhancement in virtual environments. In terms of training enhancement, shared control and virtual fixtures were no better than practice in an unassisted mode. For manual control tasks, such as the one described in this paper, shared control is beneficial for performance enhancement, but may not be viable for enhancing training effectiveness.
[DOI: 10.1115/1.2168160]

1 Introduction

The addition of haptic feedback to virtual environment simulations and telerobotic systems is known to provide benefits over visual-only displays, such as reduced learning times, improved task performance quality, increased dexterity, and increased feelings of realism and presence [1–7]. Haptic feedback in virtual environments also enables a wider range of applications, including manipulation and assembly tasks where force cues are necessary and medical applications, such as training for palpation, needle insertion, minimally invasive surgery, and rehabilitation [8]. To date, however, virtual environment systems with haptic feedback have not fully exploited the capabilities of a haptic display device. The simulators primarily focus on obeying the physical laws that govern such systems in order to recreate realistic environments for assembly, surgery, flight, and other procedures. The addition of haptic virtual cues and active assistance from the device, not realized in the physical world, can dramatically increase the amount of information that can be conveyed to a user, ideally improving performance in the virtual environment or improving the effectiveness of a haptic training or rehabilitation system.

This paper presents a shared-control interaction paradigm applied to haptic interface systems, with experimental data from two user studies. In terms of performance enhancement, shared control between a human and a robotic interface can boost performance because the robot can control low-level functions (reducing oscillation or tremor, force management, obstacle avoidance, or control of orientation) while the human operator maintains high-level control, such as path planning and position control. The area of teleoperation has seen much activity in shared control as researchers implement these partitioning techniques. This approach to human-robot interaction exploits the pros of each system for the betterment of overall performance. In addition to performance enhancement during haptic virtual environment interactions, the shared-control paradigm has implications for improving training effectiveness by reducing learning times and improving retention of manual skills, and has benefits over existing interaction paradigms, which will be discussed in Sec. 2.

The remainder of the paper is organized as follows: Section 2 provides a thorough literature review, which is necessary to define the proposed shared-control approach in relation to prior approaches for performance enhancement and training in haptic virtual environments. Section 3 describes the dynamic manual task used for experimental validation of shared control for performance enhancement and training. Section 4 describes the assistance modes, including a virtual fixture approach and the shared-control approach, with an explanation of the effect of these modes on performance of the task. Section 5 describes the experiments in detail. Experiment 1 studies the effect of shared control on task performance, while experiment 2 looks at shared control for training by measuring subject performance of a task over time. Section 6 presents the experimental results, which are discussed in Sec. 7. Finally, conclusions and implications for future work are given in Sec. 8.

2 Review of Performance Enhancement Using Haptic Devices

Most efforts to incorporate haptic feedback for performance enhancement and training have focused on three approaches. In the first approach, the haptic device is used to produce virtual force fields or fixtures in order to show the human user where not

Contributed by the Dynamic Systems, Measurement, and Control Division of ASME for publication in the JOURNAL OF DYNAMIC SYSTEMS, MEASUREMENT, AND CONTROL. Manuscript received April 2, 2005; final manuscript received November 30, 2005. Review conducted by Suhada Jayasuriya.

to go in the virtual environment. These passive guides are only perceived by the user when forbidden regions of the workspace are explored. The second approach, primarily for the purpose of training with haptic devices, is a record-and-replay strategy. Here, an expert's interactions with a virtual environment are recorded, and then in subsequent trails, a novice user feels the desired motion and then attempts to mimic it without haptic feedback. The third approach, shared control, has been used primarily for performance enhancement. The application of shared control between humans and robots for training is a relatively recent research thrust, and the relevant prior art will be discussed. In summary, this section will review advances in the related areas of passive haptic assistance for performance enhancement and training, record-and-replay strategies, and shared control for performance enhancement in both virtual and remote environments, and will lay a framework for haptic assistance for training applications.

2.1 Passive Assistance for Performance Enhancement. The premise of this paper is that virtual-force cues combined with haptic feedback due to interactions with a virtual environment can provide additional benefit to the human operator in terms of performance enhancement or even to improve training effectiveness. Some prior work has already addressed the addition of physically nonrealizable virtual cues in haptic simulations. For example, Rosenberg introduced the notion of virtual fixtures as perceptual overlays for enhanced operator performance in telemanipulation tasks [9]. In Rosenberg's work, the virtual fixture was a simulated surface that prevented hand motion beyond the surface thus guiding the user to follow the trajectory defined by the extents of the virtual fixtures. In this case, the fixtures were present throughout the duration of the teleoperation task. Because the user would always interact with the remote environment via the force-feedback interface, the virtual fixtures could be permanently overlaid on the display. A similar application is the use of virtual fixtures in a refueling task [10].

In the case of haptics for training applications, however, the trainee must learn how to interact with the simulated environment with the end goal being unassisted interaction with a real environment. Therefore, virtual fixtures may be helpful in training, but Rosenberg's fixtures were either "on" or "off" and appeared in fixed locations with fixed parameters for force-output calculation during interaction. An on-or-off approach to the use of virtual fixtures in training may improve performance during training, but performance in an unassisted virtual task or in an equivalent real-world task may be, in fact, worse since the subject has learned to rely on the presence of the virtual fixtures in the virtual training environment. This phenomenon has been seen when augmented feedback, such as computer enhancement of the environment, has been added to a graphics-only virtual environment [11].

Virtual fixtures were implemented for robot-assisted orthopedic knee surgery at Imperial College, London [12,13]. During preoperative surgical planning, shape and orientation of planned regions of resection are computed with the aid of a three-dimensional (3D) model. During the surgical procedure, the surgeon guides the robot holding the motor of a rotary cutter that is attached to the end of the robot arm. The system, dubbed Acrobot, provides assistance in achieving planned cuts with accuracy, through the use of force-control techniques. Once in the force-control mode, the behavior of the robot is determined by its relative position to the constraint boundary. Inside the hard constraint region, the stiffness displayed to the surgeon decreases with distance from the boundary and is zero inside a smaller envelope. A simple look-up table is used to determine the region in which the robot end-effector lies. Another example of virtual fixtures was presented by Bettini and others for use as manipulation assistance for interactive tasks between humans and robots [14]. They implemented soft constraints in order to allow the user to maintain more control of device positioning. The system used vision to sense a desired path in a plane, and the robot encouraged motion toward and along the path using a direction-based control law. As with virtual fixtures for teleoperated tasks, the virtual fixtures implemented with the Acrobot and Steady-Hand robotic systems can remain present for all interactions with the system since the target tasks will be performed using the robot for assisted manipulation. The results presented in this paper for performance enhancement with virtual fixtures for haptic assistance complement these prior findings and show the applicability of virtual fixtures for performance enhancement to higher-order manipulation tasks than the zero-order (position-tracking) or first-order (velocity-control) tasks summarized here. Active assistance in the form of shared control is therefore compared to virtual fixtures for assistance to determine if such interactions provide added gains in terms of human performance for second-order manual control tasks, since virtual fixtures provide only a position constraint.

2.2 Passive Assistance for Training: Record/Replay. Several researchers have sought to use haptic devices as virtual teachers. In these instances, desired trajectories or interaction forces are programed or recorded during an initial trial. Then, the end user begins a training session where this desired trajectory or force information is played back. Gillespie et al. used this method to display the optimal trajectory for balancing an inverted pendulum attached to a one-degree-of-freedom cart in the shortest possible amount of time [15]. Subjects felt the optimal trajectory of the cart and then attempted to play back the motion after the teaching portion of the simulation was completed. The subjects remained passive during the teaching phase of the experiment. Likewise, the virtual teacher was not active during the trainee's execution of the pendulum-balancing task.

This strategy has also been implemented in robotic rehabilitation systems. Burgar and his colleagues have studied the ability of a device mirror-image movement enabler (MIME) to assist limb movements and facilitate recovery of motor function in subjects with chronic hemiparesis due to stroke [16]. MIME incorporates an industrial robot and operates in three unilateral modes and one bimanual mode. In unilateral operation, passive, active-assisted, and guided movements against a resistance are possible. The MIME system's novel bimanual mode allows subjects at any impairment level to practice and complete mirror-image bimanual movements, where the unimpaired limb defines the trajectory along which the robot will carry the impaired limb. Each of these modes is equivalent to a record/replay mode where desired trajectories are commanded to the robot, and the robot carries the patient's limb through those trajectories.

Other work has attempted to record and, subsequently, play back both position and force information to the trainee [17–20]. Although these attempts succeed in displaying both position and force information, all were based on a record and replay strategy. With this approach, the subject passively felt the desired interaction modes and then, without artificial cues, tried to recreate the virtual teacher's methods. Specific outcomes for work by Huang et al. included improved human performance of a ball-and-beam–balancing task with haptic feedback during the demonstration sessions [20]. Although these methods may be successful for training, they do not take advantage of the ability of a haptic device to display corrective feedback during interactions. In other words, although the record/replay strategy is capable of demonstrating higher-order control schemes than the virtual-fixture approach, the subject remains passive during the replay mode, and when the subject is actively moving through the device's workspace, there is no corrective feedback. The proposed use of shared control for training in a haptic environment capitalizes on this capability by displaying the dynamics of the virtual system together with forces that demonstrate preferred execution techniques.

2.3 Shared Control: A Paradigm for Performance Enhancement and Training. As of late, there are a number of examples of shared control for performance enhancement in virtual environments and in teleoperation systems. Steele and Gillespie implemented shared control in a driving simulator for a path-following task [21]. They were able to reduce visual demand and improve path-following performance with the shared-control approach. Another group used shared control to suppress the swing of a crane via haptic feedback to the operator [22]. The haptic operational assistance, as it was called, employed pager motors to notify the operator of desired motions. Haptic feedback has been used for minimally invasive surgery and training, in order to improve human performance [23]. This group employed two modes of feedback for positioning and steering assistance. In each of these cases, the shared control provided feedback on zero-order (position) or first-order (velocity) manual control tasks. For the system described in this paper, the effect of shared control to improve performance of a second-order manual control task is investigated.

In teleoperation applications, shared control as a means of improving performance has typically been handled by a partitioning approach, where the human controls high-level decision making and the robotic system controls low-level operations, such as obstacle avoidance and force management [24–32]. Although this has strong implications for performance enhancement by reducing cognitive load on the operator, it is unclear if such a hierarchical approach to shared-control implementation will have merit for skill-transfer applications, where the operator must ultimately learn to control all aspects of the task. Consequently, the authors have modeled their shared-control architecture to support energy and command flow during task completion as proposed by Payandeh [33], who classifies hierarchical approaches for teleoperation more broadly as collaboration and states that shared control is a specific subset of collaboration that involves simultaneous control of a common process via energy and command flow. In this context, shared control is presented as a concept for a robotic surgical tool [33]. In the system, the human user and shared controller are in direct physical communication and power exchange. The shared controller can, for example, control the maximum transmission force regardless of the input commands from the user. As in prior examples, such a technique has clear benefit for performance enhancement, but falls short when skill transfer is the desired outcome of the interaction between the human user, robotic device, and shared controller. This is because such constraints on force magnitude do not necessarily demonstrate to the trainee the proper method of interacting with the controlled system. The shared controller, as implemented in the training experiments described herein, ensures that the subject is displayed active haptic cues that demonstrate dynamically how to control the second-order system, in addition to the full dynamics of the virtual mechanical system.

The work summarized above, together with prior work by the authors, has motivated a more thorough study of shared control as it applies to performance enhancement and skill transfer for second-order manual control tasks in haptic virtual environments. Initially, the authors found that shared control in a haptic virtual environment improved performance of a dynamic targeting task, where subjects interacted with a sprung mass and performed a hand-eye coordination task [34]. Passive assistance in the form of penalty-based, virtual-fixture-type feedback, and active assistance, now termed shared control, were implemented for performance enhancement. Shared control was formulated such that the user was directed to control the two-degree-of-freedom haptic device to suppress swing of an unactuated degree of freedom of the virtual dynamic system. As a follow-up study, the authors studied the passive-assistance and shared-control modes for training of the task [35]. These studies, the first to directly compare passive virtual fixtures and shared control, were preliminary, with insufficient subjects for rigorous statistical analysis. However, the results indicated that shared control could be at least as good as virtual fixture methods for performance enhancement and training, and had the potential to outperform virtual fixtures in the case of dynamic interaction tasks when the controlled system was of higher order than prior position or velocity control tasks.

Table 1 System parameters of the two-mass spring damper system

m_1 (kg)	m_2 (kg)	k (N/m)	b (Ns/m)
0	5	100	3

This paper presents results for a similar study to directly compare passive virtual fixtures and shared control for performance enhancement and training for a second-order manual control task. Specifically, this paper presents experimental results with a larger number of subjects than the prior published pilot studies and includes thorough statistical analysis of the experimental results. Two experiments are presented—one compares performance of a dynamic target-hitting task in each of three modes (no assistance, virtual fixtures for assistance, and shared control for assistance). Subsequently, an experiment is presented that monitors performance of the task over time while subjects train in one of these interaction modes. The findings indicate that the shared-control approach is equally as good as virtual fixtures for performance enhancement for the second-order manual control task. Indeed, shared control may prove more effective than virtual fixtures due to the more general structure and implementation as compared to virtual fixtures. For training, there are no significant differences in group performance for the assistance-interaction modes applied during training (no assistance, virtual fixtures, and shared control) when unassisted performance was compared from session to session. There are some indications that shared control may be better than virtual fixtures for training enhancement and that both assistance modes result in better retention of skill from one training session to the next, compared to practice without assistance, but such claims are only loosely supported by the experimental findings.

3 System and Task Description

To evaluate the effectiveness of haptic assistance on performance enhancement and training, a system with two point masses connected by a spring and a damper in parallel was used. This two-mass system has four degrees of freedom (DOF), namely, the x and y motion of both m_1 and m_2. However, subjects can only control directly the x and y movement of mass m_1 via the haptic 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. Additionally, the task qualifies as a second-order manual control task, since the system contains two integrations between control input and plant output [36].

Such a system is well suited for an experimental study of human performance enhancement and training with haptic assistance because the motions are sufficiently complex to control, and because reflection of force feedback generated by the interactions of the two masses connected by the spring damper is necessary for the human to accurately control motion of the system. Also, haptic feedback is considered to be important for performance and learning of dynamic tasks [23]. Therefore, we can examine the forces of interaction due to the system's inherent dynamics, and those additional forces that we overlay on the environment for assistance due to passive virtual fixtures or the shared controller. Table 1 lists the parameters that govern the dynamic response of this system.

3.1 Hardware. An Impulse Engine 2000 joystick from Immersion Inc., shown in Fig. 1 was used to provide a high-fidelity haptic simulation of the two-mass system. The Impulse Engine has two degrees of freedom and a workspace of 6 in. $\times$ 6 in. The device exhibits low backdrive friction (<0.14 N) and a high sensor resolution (0.0008 in.). All simulations ran at the sampling

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

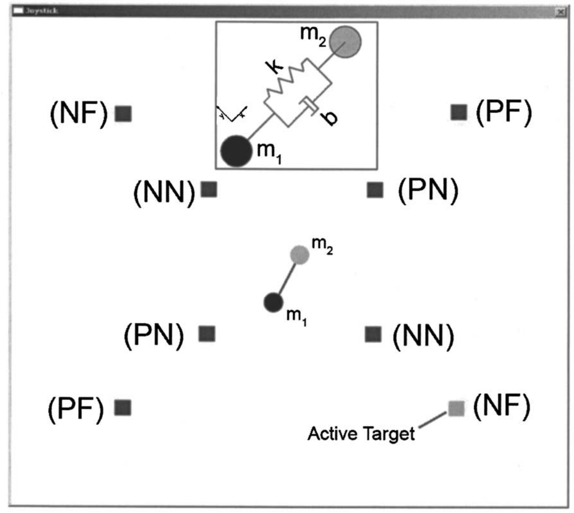

Fig. 2 Graphical display of tapping experiment. Subjects control location of m_1 in order to cause m_2 to hit the desired target. Targets appear in pairs (NF: negative slope, far; NN: negative slope, near; PF: positive slope, far; PN: positive slope, near). Inset shows virtual underactuated system. The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

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 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. After preliminary experiments with the authors as test subjects, the values for m_2, k, and b were chosen to be 5 kg, 100 N/m, and 3 Ns/m, respectively, to ensure the system to be easily controllable.

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 relatively low inertia.

3.2 Task. A target-hitting task is used to study human 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 four target pairs that are utilized in the experiments. They are referred to as follows: positive slope near (PN), positive slope far (PF), negative slope near (NN), and negative slope far (NF). 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 [34]. 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. Performance of the task was measured by time between target hits.

4 Haptic-Assistance Modes

The goal of the experiments was to investigate the efficacy of different haptic-assistance modes for performance enhancement and training for the task described above. The haptic assistance here refers to the additional forces displayed to the subjects via the haptic joystick, and the method selected to display assistance forces constitutes the independent variable of the human subject experiment. Specifically, the authors sought to determine if the shared control paradigm was more effective than overlaid perceptual cues, such as virtual fixtures for improving performance or training of the second-order manual control task. Virtual fixtures, commonly used for performance enhancement in teleoperator systems, were implemented as a passive type of haptic assistance of zero order. In contrast, the shared-control paradigm for haptic assistance represents active intervention and takes account of the second-order behavior of the controlled system.

Virtual fixtures are a penalty-based form of haptic assistance dependent on the subject's position within the virtual environment. When the subject moves to a forbidden region of the workspace, or veers off the intended path, corrective feedback in the form of soft virtual walls serve to push the haptic device, held by the user, back to an acceptable position. In contrast, shared control is an active interaction that depends on the dynamic system that the subject is controlling. For the task described in this paper, assistance of the virtual fixture type constrains the position of the user, while assistance in the form of shared control applies controlled 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. Shared-controller architectures could take many forms. For example, the controller could be implemented to reduce the apparent order of the controlled system from the perspective of the user. Another option, as implemented in this work, is to have the shared controller reduce the difficulty of the task by altering the dynamics of the controlled system (e.g., suppressing the motion of the disk normal to the target axis). In either case, the output of the shared controller, and, in turn, the assistance forces, are a function of all of the state variables of the virtual or remote dynamic system, not just the independent variables controlled by the user's motion input to the haptic device. For simple tasks, virtual fixtures are likely to be sufficient for

Table 2 Control parameters

Virtual fixture		Shared control		
k_{wall}=22.8 N/m	b_{wall}=0.57 Ns/m	λ=1	k_p=70	k_v=1

performance enhancement and even for training, as the results of this paper indicate. For more complex tasks, such as those requiring manual control of higher-order systems (e.g., cranes, flexible robots, or multi-degree-of-freedom manipulators), the shared-control architecture provides additional flexibility in implementation.

The three interaction modes used in the performance-enhancement and training experiments (no assistance, virtual fixtures, and shared control) are described in detail in the following sections, with specific control parameters for each mode summarized in Table 2.

4.1 Dynamics of the Virtual Sprung-Mass System. For the purpose of following discussion, the x-axis is considered to be aligned with the positive-slope target pair, and the y-axis is perpendicular to x with the positive direction forming a right-handed coordinate system. The dynamics of the spring-mass system in our study 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 are the x and y positions of masses m_1 and m_2, respectively. F_{kx} and F_{ky} denote the x and y components of the forces arising from the spring and the damper (see Table 1 for parameter values). Explicitly, they are calculated according to

$$F_{kx} = k(x_2 - x_1) + b(\dot{x}_2 - \dot{x}_1) \tag{5}$$

$$F_{ky} = k(y_2 - y_1) + b(\dot{y}_2 - \dot{y}_1) \tag{6}$$

F_x and F_y are the external forces exerted on the mass m_1 through actuators of the haptic device.

4.2 No Assistance. As evident from the name, this mode serves as the control set and no haptic assistance was provided. 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 virtual fixture and shared-control cases, subjects felt the forces due to both the internal dynamics of the system and the augmented forces intended to assist in task completion.

4.3 Virtual Fixture Assistance. In the virtual fixture assistance mode, a pair of virtual walls, modeled as a spring and damper in parallel, applies forces on the subject's hand in case of deviation from the reference motion of mass m_1. For instance, if the target pair is aligned with the x-axis, then motion of the joystick, which controls the motion of m_1 in the positive y direction will result in a force applied to the joystick in the negative y direction. In this assistance mode, virtual walls were used to encourage users in a passive manner to move mass m_1 along the axis between targets, under the assumption that such motion of m_1 would tend to cause m_2 to move generally along the same path. The virtual wall generated forces proportional to the error in position of mass m_1 measured in the direction normal to the axis of the target pair, and proportional to the velocity of m_1 in the normal direction. The force for one virtual wall for a target pair aligned with the x-axis is calculated according to

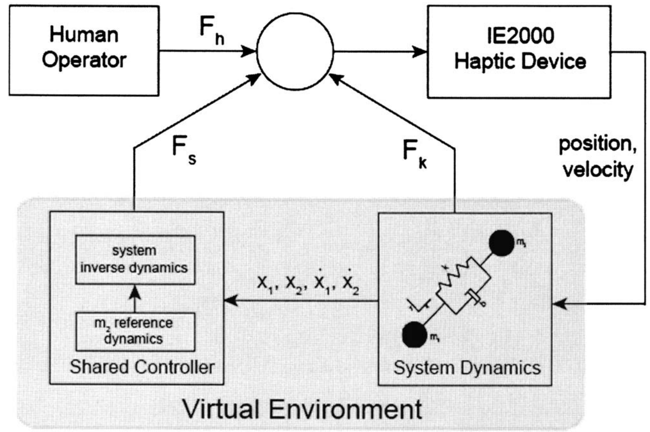

Fig. 3 Block diagram of shared-controller architecture. The position and velocity of the joystick are mapped directly to that of m_1. The shared controller computes the force F_s (Eqs. (10) and (11)) that should act on m_1 in order to follow the reference dynamics for m_2, as defined by Eqs. (8) and (9). The force F_s is reflected back through the haptic device in addition to the forces arising from system dynamics, F_k (Eqs. (5) and (6)).

$$F_{py} = k_{\mathrm{wall}}(y_1 - y_{\mathrm{wall}}) + b_{\mathrm{wall}}\dot{y}_1 \tag{7}$$

These forces were subsequently displayed to the user via the 2-DOF haptic joystick. The virtual wall parameters were chosen to be k_{wall}=22.8 N/m and b_{wall}=0.57 Ns/m.

4.4 Shared-Control Assistance. In the shared-control assistance mode, the motion of m_2 instead of m_1 as with virtual fixture assistance, is constrained along the active 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. 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 Eqs. (1)–(4). The implementation of the shared-control mode is represented in block diagram form in Fig. 3.

A simple feedback controller can be implemented for position control of mass m_2 without explicitly deriving F_{kx} and F_{ky} in Eqs. (1)–(4). Consider an active target pair aligned along the x-axis. The desired controlled dynamics for m_2 in the y direction are defined as

$$\ddot{x}_2 + 2\lambda\dot{x}_2 + \lambda^2 x_2 = -K_v\dot{x}_2 \tag{8}$$

$$\ddot{y}_2 + 2\lambda\dot{y}_2 + \lambda^2 y_2 = -K_p y_2 - K_v\dot{y}_2 \tag{9}$$

where K_p and K_v are control gains.

F_{sx} and F_{sy}, the forces to be displayed due to the shared controller alone, can be solved by eliminating $\ddot{x}_2$ and $\ddot{y}_2$ from the set of equations (8) and (9) using Eqs. (1)–(4), as shown in

$$F_{sx} = m_1\ddot{x}_1 - m_2[(K_v + 2\lambda)\dot{x}_2 + \lambda^2 x_2] \tag{10}$$

$$F_{sy} = m_1\ddot{y}_1 - m_2[(K_v + 2\lambda)\dot{y}_2 + (K_p + \lambda^2)y_2] \tag{11}$$

Equations (10) and (11) represent the proposed shared-control methodology. Similar analysis can be used to derive the shared-controller assistance forces for the other target pair. This approach to shared control is unique in that it is nonhierarchical, since the forces displayed to the joystick include the system dynamics of the virtual two mass system and the output of the automatic controller, which is acting to suppress the off-axis swing of m_2. The inertia term corresponding to the inertia of mass m_1 (the joystick in this case) was considered negligible, due to the high back derivability of the IE2000. The parameter set for the shared-control assistance mode is: λ=1, k_p=70, and k_v=1.

4.5 Comparison of Assistance Modes. This section seeks to illustrate the effect of each assistance mode described above. First, the typical force profiles for interactions in each mode are presented and the contribution of the virtual environment forces and assistance forces to the total commanded force, sensed by the user, are contrasted. Then, traces are presented to show the resultant effects of the assistance modes on the trajectories of masses m_1 and m_2.

4.5.1 Force Profiles. Figure 4 shows a force profile for interaction with a target pair in the no-assistance mode. Recall that the x- and y-axes are aligned with the target axes, shown in Fig. 2. In Fig. 4(*a*) it is clear that force feedback is primarily in the x direction, since the user is attempting to cause m_2 to hit targets aligned with the x-axis. F_x and F_y are due entirely to the forces that arise from the user's control of the motion of m_1 via the joystick, and the resultant dynamics due to the spring and damper system between m_1 and m_2. Figure 4(*b*) shows a representative force profile for interaction with a target pair in the presence of the virtual fixture assistance mode. Note significant forces in the y direction despite the target-pair alignment being along the x-axis. The magnitude of F_y is due to contributions from the spring-mass-damper system forces (F_{ky}) and due to the virtual fixture feedback (F_{py}). Since the target pair for this trial is aligned along x, motion in the y direction leads to force feedback from the virtual fixture. This is the primary component of F_y, the total y-axis force displayed with the haptic device. A force profile for a representative interaction with a target pair oriented along the x-axis, in the shared-control assistance mode, is shown in Fig. 4(*c*). In the plot, F_x and F_y are the total forces displayed to the user, the sum of the F_k and F_s components described in (5), (6), (10), and (11). F_{kx} and F_{ky} are components of the force due to excitation of the spring-mass-damper system, and F_{sy} is the component of reflected force due to the shared controller. F_{sy} is calculated, based on the above equations, such that the y-axis motion of m_2 is suppressed.

It should be emphasized that both the virtual fixture and shared-control assistance modes only affect the motion of the constrained mass in the direction perpendicular to the orientation of the active target pair. If the joystick is released by the user, the handle will move such that the motion of either m_1 or m_2 is constrained in this perpendicular direction. However, the control algorithms will not cause m_2 to be drawn toward the active target. Therefore, the subject is required to initiate the corner-to-corner motion of the joystick necessary to cause successful performance of the task.

4.5.2 Traces. To better articulate the resulting motion of the underactuated system due to the presence of each assistance mode, traces of the paths of m_1 and m_2 were recorded for a single trial in each mode and are presented in Fig. 5. The gray lines are the traces of m_2 whereas the black lines trace the motion of mass m_1. In the unassisted mode, the subject is free to move about the workspace. In the virtual fixture assistance mode, the movement of m_1 is constrained to be along an axis between the target pair due to the presence of the virtual walls. In the shared-control assistance mode, the motion of the mass m_2 is restricted to move along the target axis due to the action of the shared controller.

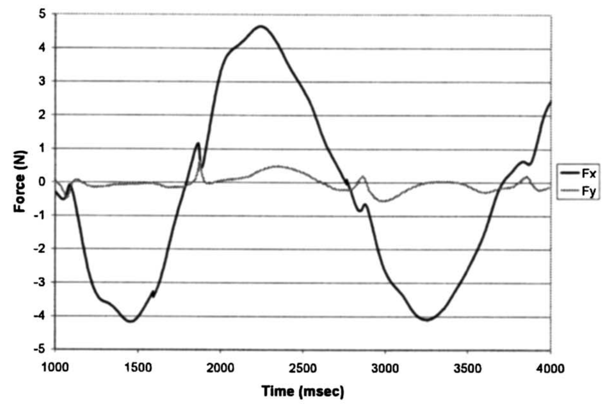

(a)

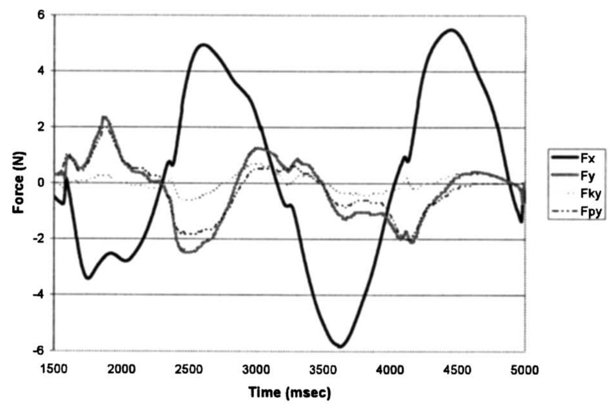

(b)

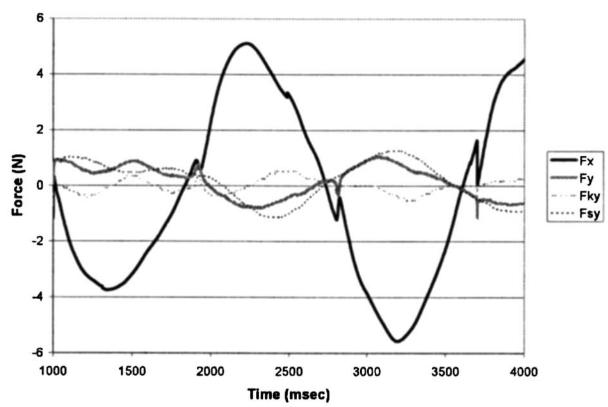

(c)

Fig. 4 Force profiles for interactions with a target pair aligned with the x-axis with (*a*) no assistance, (*b*) virtual fixtures for assistance, and (*c*) shared control for assistance. In (*a*), the total commanded force felt by the user, F_x and F_y, is due entirely to the forces that arise from the dynamics of the virtual spring-mass system, and due to the target alignment, the predominant forces are in the x direction. In (*b*), note the existence of significant forces in the y direction despite target-pair alignment equivalent to case (*a*). The magnitude of F_y is due to contributions from the spring-mass system forces (F_{ky}) and the virtual fixture feedback (F_{py}). Since the target pair for this trial is aligned in x, motion in the y direction, leads to force feedback from the virtual fixture. This is the primary component of F_y, the total y-axis force displayed with the haptic device. In case (*c*), F_{ky} is the force due to excitation of the spring-mass-damper system, while F_{sy} is the force due to the shared controller. F_{sy} is calculated such that the y-axis motion of m_2 is suppressed. Similar to (*b*) F_y is significant, but the forces are a function of all of the system's state variables, not just the position of m_1.

5 Experiment Details

Fifteen subjects, four females and 11 males, participated in the experiments. Two of the subjects were left-handed, and only a few subjects had prior experience with the haptic interface or the task under study. After a practice session to orient subjects with the haptic interface device and the task, the subjects began with a single performance session and then completed nine training sessions over four weeks.

5.1 Experiment 1: Performance. All subjects complete an identical performance-assessment session to evaluate the effectiveness of each assistance mode on performance enhancement.

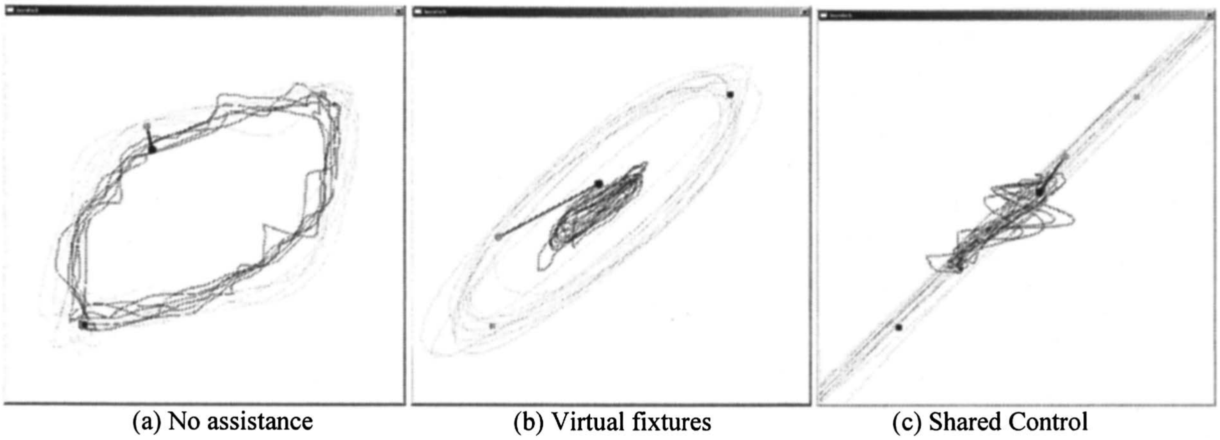

Fig. 5 Traces show representative paths of m_1 (black) and m_2 (gray) in the workspace in each mode of interaction. The no assistance mode (*a*) constrains neither the motion of m_1 nor the motion of m_2. In the virtual fixture assistance mode (*b*), the motion of m_1 is constrained to lie along a straight-line path between the target pair. Subjects experience a resistive force proportional to the distance deviated from this line. The motion of m_2 remains unconstrained. In the shared-control assistance mode (*c*), the motion of m_2 is controlled such that any motion orthogonal to the path between the targets is suppressed, therefore, making it relatively easy for subjects to maintain motion of m_2 along the target axis. The motion of m_1 in this mode is not directly constrained.

The performance session consists of 120 trials, where a trial is defined as a set of 13 sequential target hits with a particular combination of assistance mode (three levels—no assistance, virtual fixture assistance, and shared-control assistance), target orientation (two levels—positive and negative slope), and target distance (two levels—near and far), for a total of 12 possible combinations. During the performance session, each combination of assistance mode, target orientation, and target distance is presented to the subject ten times, and the order of presentation is randomly assigned throughout the experiments. All users train with the same system and have no prior knowledge of the task. Performance was measured as the time between target hits. The intertap time is related to the natural frequency of the system, and both are dependent on the mass-spring-damper system parameters selected for the task. Although intertap time provides a useful quantitative measure of performance, it should be noted that this measure is specific for this given task and parameter set. Throughout the experiment and across all subjects, the required task and the dynamics of the system remain unchanged. Thus, measurement of performance in terms of intertap provides a metric for comparative analysis of learning and user performance.

5.2 Experiment 2: Training. A training experiment is conducted after the performance experiment to investigate the influence of virtual fixture and shared-control haptic assistance modes on human learning of the dynamic task. This experiment is termed training rather than an adaptation study since subjects are asked to learn a pattern of motion that indirectly controls the position of m_2 via motion of m_1. This study can be categorized as learning of a new motor skill, whereas adaptation is a process that takes place under perturbation force fields. In such adaptation studies, after effects are the intended result of perturbation, and such effects wash away quickly. In the training experiments, performance both with and without the haptic assistance forces is monitored over several weeks, and human motor performance is notably altered in this time. Subjects do not revert back to their preexposure manipulation paths, but learn the motor skill as demonstrated by the assistance forces or by exploring their own manipulation strategies.

During the course of the training experiment, subjects were divided into three groups by assistance mode, namely, no assistance, virtual fixture assistance, and shared-control assistance. Subjects completed nine sessions, with each training session consisting of 40 trials (two levels of orientation, two levels of distance, and ten repetitions of presentation). 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 session. For each baseline test, subjects completed 20 trials (two levels of orientation, two levels of distance, and five repetitions of presentation), all in the no-assistance mode. A training session and its corresponding two 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.

6 Results

Performance-assessment results for the three interaction modes are presented in Fig. 6. Standard errors for the results are shown with error bars. The results are as the reader might expect, with haptic assistance, both from the virtual fixture approach and from

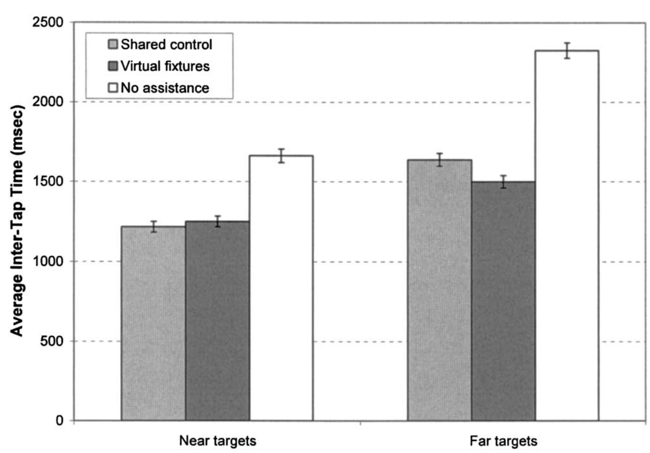

Fig. 6 Experiment 1: Average time between target hits for performance assessment experiment. Shared control and assistance using virtual fixtures both provide significant and similar improvement in user performance as compared to the no-assistance case. A significant interaction was found between distance and assistance modes. The results for positive and negative sloped target pairs are combined in this figure.

the shared-control approach, causing a marked improvement in performance of the task as seen by a decrease in the average time between target hits. Performance in the virtual fixture and shared-control assistance modes is significantly better than in the unassisted mode. Performance gains are larger for the case where targets are father apart. 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 experimentwise error rate, whereas independent t tests do not. Analysis of variance results show that assistance mode ($P < 0.0001$) and target distance ($P < 0.0001$) are significant effects. F-test contrasts of the experiment variables indicate that shared control and no assistance are significantly different ($P < 0.0001$), as are no assistance and virtual fixture assistance ($P < 0.0001$). Shared-control and virtual fixture assistance do not produce significantly different results. Two-way interactions were also analyzed, with a significant interaction between assistance mode and target distance ($P < 0.0001$).

Training experimental results are presented in Fig. 7. The average intertap times for the before and after baseline measurements are given for each day of the training. Note that the baseline tests were conducted in the absence of any assistance mode, while the subjects trained with one of the three modes (no assistance, virtual fixture assistance, or shared-control assistance). At the conclusion of training, subjects in each group seem to converge to the same performance level. Subjects in the no-assistance group performed better in initial sessions, which could be due to subjects in this group being inherently better at the dynamic targeting task from the beginning. The virtual fixture and shared-control groups reach the same performance levels as the control group by the fourth or fifth day of training. If initial performance of the control group had been comparable to that of the virtual fixture or shared-control groups, perhaps more significant conclusions could be drawn regarding the effects of haptic assistance for performance enhancement.

For statistical analysis, the following were treated as factors: session (one through nine), assistance mode during training (none, virtual fixtures, or shared control), target slope (positive or negative), target distance (near or far), and baseline measurement (before or after). Session ($P < 0.0001$), assistance mode ($P = 0.0101$), target distance ($P < 0.0001$), and baseline measurement ($P = 0.0025$) were significant. Two-way interactions were analyzed, with significant interactions for the following combinations: session and assistance mode ($P < 0.0001$), session and target distance ($P < 0.0001$), session and baseline ($P = 0.0081$), assistance mode and target slope ($P = 0.0041$), and assistance mode and target distance ($P = 0.0005$). For F-test contrasts between experiment variables, performance of the shared-control versus no-assistance groups was not significant. Similarly, performance of the shared-control versus virtual fixtures groups was not significant. However, performance of the no-assistance group was significantly different from the virtual fixture group ($P = 0.0030$).

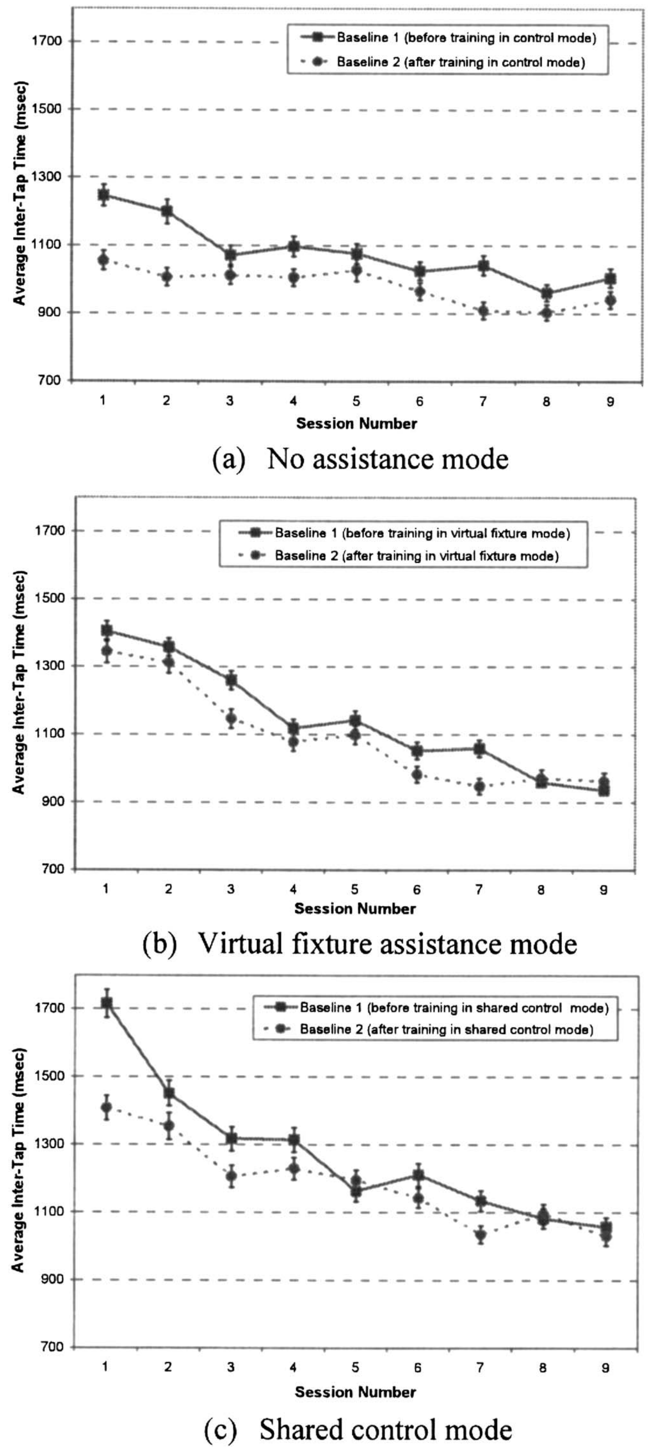

Fig. 7 Experiment 2: Baseline performance measures across training sessions. Subjects' performance in an unassisted mode was measured before and after a training session in one of three assistance modes ((*a*) no assistance, (*b*) virtual fixture, or (*c*) shared control). Results for near and far targets are combined in this presentation. The lower bound on performance is approximately 700 ms, limited by the natural frequency of the virtual mass-spring-damper system.

7 Discussion

The results from the performance experiment with and without haptic assistance support prior findings, where passive assistance in the form of virtual fixtures gives rise to better performance than an unassisted mode. The authors hypothesized that the shared-control assistance mode would result in better performance than the virtual fixture assistance mode, since the shared controller assisted the user in completion of the second-order manual control task. This was not the case, however. This could be because the task required motion along a linear path for best results, and the virtual fixture mode was able to display this motion passively to the human. As expected, the slope of the target pair was not a significant factor in the experiment, and therefore, for the results presented in Fig. 6, positive and negative slope pairs are combined. Distance between target pairs was significant, which is because the greater distance between targets provides an opportunity for undesired motions of the underactuated system to be commanded. The significant interaction between assistance mode and target distance implies that when the task is more difficult (target distance is greater), the assistance modes give rise to larger performance gains. This conclusion is based on two observations. First, the average intertap times with virtual fixtures and with shared control are better than the no-assistance case, and there is a greater net difference in performance between no assistance and virtual fixtures and between no assistance and shared control when the targets are spaced farther apart. Second, the statistical analysis shows that the results for performance with virtual fixtures and with shared control are not significantly different. For these reasons, the authors conclude that haptic assistance of either type is more beneficial than no assistance when the targets are farther apart (i.e., the task is more difficult).

The conclusion to be drawn from the performance-assessment experiment is that shared control gives rise to the same performance gains as virtual fixtures for assistance. If performance enhancement in a virtual or teleoperation system is desired, virtual fixtures are sufficient, and a shared-control architecture is not required. This is advantageous because of the passive nature of the virtual fixture approach, where assistance feedback is displayed by penalty-based methods, rather than the shared control case where feedback does not directly depend on the motion input of the human operator. It is likely, however, that if the task was more complex, for example, if the targets changed location or if the dynamics of the system were changed during a trial, the shared-control approach may prove more beneficial than virtual fixtures. There are a number of reasons the authors feel that more complex tasks might benefit from shared control. First, results from the training experiments show a saturation of performance prior to the end of training. The authors feel that a more complex task, which takes longer to learn, may provide more data for analysis and may elucidate additional differences between the assistance approaches, namely, virtual fixtures and shared control. Specifically comparing virtual fixtures and shared control for assistance, the authors contend that shared control, by its active nature, has the potential to simplify the dynamics of a system with which a user interacts. It is well known in the human motor control literature that performance degrades for higher-order tasks. Jagacinski and Flach [36] state that task of controlling second-order systems (e.g., vehicular control, remote manipulation, or controlling systems, such as the sprung mass employed here) is difficult, such acceleration control can be mastered by most people with practice. Third- and higher-order control, common in the fields of aviation, are quite difficult for people to control. That being said, mastery of such tasks can be achieved, as is seen by skilled pilots who can become proficient with proper training and feedback displays [36]. Shared control offers added flexibility for implementation as haptic assistance when compared to virtual fixtures. For example, shared control can be used in a demonstrative phase, and the feedback can be a function of user performance. In other words, the complexity of the task can be gradually increased by adjusting the gains of the shared controller, adaptively, based on performance metrics. The assistance provided by the shared controller then is essentially phased out as the subject masters the task. Another possibility is to use a shared controller to effectively reduce the order of the system and isolate aspects of the task that are to be controlled by the subject, which may have benefit for performance enhancement. Other possibilities for training in haptic environments that incorporate shared control include displaying more complex dynamics instead of simpler ones, or varying system parameters in parallel with error augmentation. Virtual fixtures, by their passive nature, cannot demonstrate how to do a task, but can only provide penalty-based feedback based on the user's motion when completing the given task.

In the training experiment results shown in Fig. 7, improvement in performance is noted in same-day baseline measurement comparisons and in day-to-day comparisons throughout the training procedure. However, the expected result, that training in the shared-control mode would give rise to larger performance gains than the unassisted or virtual fixture modes, is not seen. Performance for the group that trained in the unassisted mode converges to the same level as those subjects who trained with virtual fixtures or shared control for assistance. In fact, for this group of subjects, initial performance of the control group was better than performance of the virtual fixture or shared-control groups. Since all baseline experiments were conducted in an unassisted mode, all groups had the same amount of exposure to the task for the first baseline experiment. Therefore, high scores for the first baseline experiment indicate a higher level of initial skill for this task. Future work will incorporate evaluation measures of baseline skill prior to assigning subjects to groups in order to ensure adequate distribution for post hoc analysis. All of the groups approach, but do not reach, the optimal intertap time of 700 ms determined by the system's natural frequency. The selection of a more difficult task for training comparisons may lead to stronger conclusions about the efficacy of shared control for assistance during training, as compared to no assistance (practice) or virtual fixtures for assistance. However, based on the results presented here, practice in an unassisted mode is as good as training with haptic assistance, be it due to virtual fixtures or from an automatic controller. Although performance of the task with virtual fixture or shared-control assistance gives rise to better performance, this gain does not result in improved performance among subjects who train with assistance and then are measured without assistance. One hypothesis is that the on/off approach to assistance used here, be it virtual fixtures or shared control, is not the most effective approach to training of a dynamic task in a virtual environment. This phenomenon has been seen when augmented feedback, such as computer enhancement of the environment, has been added to a graphics-only virtual environment. Todorov [11] and others noted that although some forms of augmented feedback in a graphical virtual environment were shown to enhance learning of simple movements, the performance gains achieved during training seldom transferred to the real task [31]. Long-term retention of the skill has yet to be tested across training groups.

A final representation of training results is shown in Fig. 8, which also shows the net change in performance from the end of one training session to the beginning of the next training session ($B2_i - B1_{i+1}$). The y-axis value is the raw $B2_i$ score minus the raw $B1_{i+1}$ score, averaged across subjects in each training group, where i represents the training session number. In other words, the performance at the end of a given training session is compared to the performance at the start of the following training session, with both performance measures recorded without any form of haptic assistance. Negative values indicate that the average time between hits has decreased from one session to the next, implying that the subjects' performance has improved since the last session. Positive values indicate that the average time between hits has increased (i.e., the performance has degraded since the last session). Session-to-session performance change is thus plotted for each training group (no assistance during training, virtual fixtures during training, or shared control during training). This representation shows an interesting feature of the haptic assistance modes that is not evident in the group that merely practiced the task without assistance. As seen in Fig. 8, the no-assistance group always experiences an increase in average intertap time between training sessions, except for the period between sessions 5 and 6, where performance neither degraded nor improved. For the cases of virtual fixtures and shared control for assistance during training, subjects showed a net *decrease* in performance between training sessions nearly half of the time. Upon review of Fig. 7, it is clear that all groups experience an improvement in performance from the beginning of any given training session to the end of that session and also improve from session to session if the first or second baseline scores are compared independently. The investigation of performance difference from the end of one training session to the beginning of the next training session, as shown in Fig. 8, illustrates that the no-assistance group sees performance degradation between each training session, whereas this is not the case for the

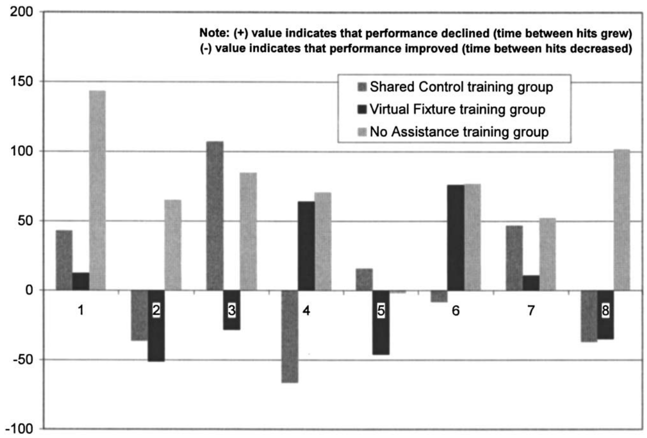

Fig. 8 Experiment 2: Regression in task performance between training sessions. Y-axis value is the raw $B2_i$ score minus the raw $B1_{i+1}$ score, averaged across subjects in each training group, where *i* represents the training session number. In other words, the performance at the end of a given training session is compared to the performance at the start of the following training session, with both performance measures recorded without any form of haptic assistance. Negative values indicate that the average time between hits has decreased from one session to the next, implying that the subjects' performance has improved since the last session. Positive values indicate that the average time between hits has increased (i.e., the performance has degraded since the last session). Only the shared-control and virtual fixture groups experience a performance improvement *between* sessions. The no-assistance group experiences degradation in performance between each training session, such that the task must be relearned before additional improvements in performance can be realized.

groups that are provided with haptic assistance, either via virtual fixtures or shared control. This fact indicates that assistance during training may have some benefits in terms of retention. Further analysis of this phenomenon will be conducted in future experiments.

8 Conclusions

This paper presents a thorough literature review of haptic assistance in virtual environments and cooperative manipulation, and shared control in teleoperation. These topics form the basis for the proposed shared-control architecture for performance enhancement and training in haptic virtual environments. The paper then presents the application of the shared-control paradigm for performance enhancement and improved training effectiveness for a second-order manual control task in a virtual environment with haptic feedback. Shared control is a form of haptic assistance for which the haptic device contributes to execution of the task via force commands from an automatic controller. Compared to haptic virtual environments that merely display the physics of the virtual system or to passive methods of haptic assistance for performance enhancement based on virtual fixtures, the shared-control approach offers a method for actively demonstrating desired motions during virtual environment interactions. Although shared control has been implemented in other studies for zero- and first-order manual control tasks, this work implements shared control for performance enhancement of a second-order task. In addition, shared control is investigated as an alternate approach to training of second-order manual tasks in haptic virtual environments.

The shared-control paradigm for assistance was compared to passive virtual fixtures and no assistance to assess performance enhancement for a dynamic manual target-hitting task. Results indicate that shared control is equally as effective as virtual fixtures at boosting performance, although shared control was expected to outperform the virtual fixture approach. This result could be due to the relatively low level of complexity of the manual task that was selected, and future work will address third- and higher-order manual control tasks. The shared-control paradigm may be of great use for performance enhancement applications since it is more general than the virtual fixtures approach. Following the performance experiment, subject performance was tracked over nine training sessions where subjects were divided into groups by training mode (no assistance, virtual fixtures for assistance, or shared control for assistance). Here, no significant gains in training performance were noted with the shared-control mode. Based on the manual control task used in this work, shared control is found to be beneficial for performance enhancement, but no better than practice for training of the manual task.

Acknowledgment

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 No. N00014-04-1-0517 and by the Rice University George R. Brown School of Engineering Undergraduate Research Program.

References

[1] Massimino, M. J., and Sheridan, T. B., 1994, "Teleoperator Performance With Varying Force and Visual Feedback," Hum. Factors, **36**(1), pp. 145–157.

[2] Richard, P., and Coiffet, P., 1995, "Human Perceptual Issues in Virtual Environments: Sensory Substitution and Information Redundancy," *Robot and Human Communication RO-MAN'95 Tokyo*, *Proceedings of 4th IEEE International Workshop*, IEEE, New York, pp. 301–306.

[3] Fabiani, L., et al., 1996, "Human Interface Using the Rutgers Master II Force Feedback Interface," *Proceedings of the IEEE 1996 Virtual Reality Annual International Symposium, Mar. 30–Apr. 3*, Santa Clara, CA, IEEE, New York, pp. 54–59.

[4] Meech, J. F., and Solomonides, A. E., 1996, "User Requirements When Interacting With Virtual Objects," *IEE Colloquium (Digest), Proc. IEE Colloquium on Virtual Reality—User Issues*, Mar. 25, IEE Paper No. 1996(068), pp. 3/1–3/3.

[5] Adams, R. J., Klowden, D., and Hannaford, B., 2001, "Virtual Traning for a Manual Assembly Task," Haptics-e, **2**(2).

[6] Williams, L. E. P., et al., 2002, "Kinesthetic and Visual Force Display for Telerobotics," *2002 IEEE International Conference on Robotics and Automation*, May 11–15, pp. 1249–1254.

[7] O'Malley, M. K., et al., 2003, "Simulated Bilateral Teleoperation of Robonaut," *AIAA Space 2003*, Long Beach, CA.

[8] Burdea, G., 1996, *Force and Touch Feedback for Virtual Reality*, Wiley, New York.

[9] Rosenberg, L. B., 1993, "Virtual Fixtures: Perceptual Tools for Telerobotic Manipulation," *1993 IEEE Annual Virtual Reality International Symposium*, Sep. 18–22, Seattle, IEEE, New York, pp. 76–82.

[10] Haanpaa, D. P., and Roston, G. P., 1997, "Advanced Haptic System for Improving Man-Machine Interfaces," Comput. Graph., **21**(4), pp. 443–449.

[11] 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), pp. 147–158.

[12] Ho, S. C., Hibberd, R. D., and Davies, B. L., 1995, "Robot Assisted Knee Surgery—Establishing a Force Control Strategy Incorporating Active Motion Constraint," IEEE Eng. Med. Biol. Mag., **14**(3), pp. 292–300.

[13] Davies, B. L., et al., 1997, "Active Compliance in Robotic Surgery—The Use of Force Control as a Dynamic Constraint," Proc. Inst. Mech. Eng., Part H: J. Eng. Med., **211**(4), pp. 285–292.

[14] Bettini, A., et al., 2002, "Vision Assisted Control for Manipulation Using Virtual Fixtures: Experiments at Macro and Micro Scales," *2002 IEEE International Conference on Robotics and Automation*, May 11–15, Washington DC.

[15] Gillespie, R. B., et al., 1998, "The Virtual Teacher," *Proceedings of the 1998 ASME International Mechanical Engineering Congress and Exposition*, Nov. 15–20, ASME, New York.

[16] Burgar, C. G., et al., 2000, "Development of Robots for Rehabilitation Therapy: The Palo Alto VA/Stanford Experience," J. Rehabil. Res. Dev., **37**(6), pp. 663–673.

[17] Yokokohji, Y., et al., 1996, "Toward Machine Mediated Training of Motor Skills—Skill Transfer From Human to Human via Virtual Environment," *Proceedings of 1996 5th IEEE International Workshop on Robot and Human Communication RO-MAN*, Nov 11–14, Tsukuba, Japan, IEEE, New York.

[18] Henmi, K., and Yoshikawa, T., 1998, "Virtual Lesson and its Application to Virtual Calligraphy System," *Proceedings of 1998 IEEE International Conference on Robotics and Automation, Part 2 (of 4)*, May 16–20, Leuven, Belgium, IEEE, New York.

[19] Kikuuwe, R., and Yoshikawa, T., 2001, "Haptic Display Device With Fingertip

Presser for Motion/Force Teaching to Human," *2001 IEEE International Conference on Robotics and Automation (ICRA)*, May 21–26, Seoul.
[20] Huang, F., Gillespie, R. B., and Kuo, A., 2002, "Haptic Feedback and Human Performance in a Dynamic Task," *Intl. Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 24–31.
[21] Steele, M., and Gillespie, R. B., 2001, "Shared Control Between Human and Machine: Using a Haptic Steering Wheel to Aid in Land Vehicle Guidance," *Human Factors and Ergonomics Society, 45th Annual Meeting*.
[22] Yoneda, M., et al., 1999, "Assistance System for Crane Operation with Haptic Display—Operational Assistance to Suppress Round Payload Swing," *IEEE International Conference on Robotics and Automation*, IEEE, New York, pp. 2924–2929.
[23] Dang, T., Annaswamy, T. M., and Srinivasan, M. A., 2001, "Development and Evaluation of an Epidural Injection Simulator With Force Feedback for Medical Training," *Medicine Meets Virtual Reality (MMVR)*.
[24] Anderson, R. J., 1996, "Autonomous, Teleoperated and Shared Control of Robot Systems," *Proceedings of IEEE International Conference on Robotics and Automation*, IEEE, New York, pp. 2025–2032.
[25] Tarn, T.-J., et al., 1996, "Task-Oriented Human and Machine Cooperation in Telerobotic Systems," Annu. Rev. Control, **20**, pp. 173–178.
[26] Williams, R. L., II, Harrison, P. W., and Soloway, D. L., 1997, "Shared Control of Multiple-Manipulator, Sensor-Based Telerobotic Systems," *Proceedings of 1997 IEEE International Conference on Robotics and Automation*, April 20–25, IEEE, New York, Vol. 2, pp. 962–967.
[27] Amorose, R. O., et al., 2000, "Robonaut: NASA's Space Humanoid," IEEE Intelligent Systems and Their Applications, **15**(4), pp. 57–63.
[28] Debus, T., et al., 2000, "Cooperative Human and Machine Perception in Teleoperated Assembly," *Experimental Robotics VII*, D. Rus and S. Singh, eds., Springer-Verlag, Berlin.
[29] Goldberg, K., et al., 2000, "Collaborative Teleoperation via the Internet," IEEE International Conference on Robotics and Automation, pp. 2019–2024.
[30] Lim, P., et al., 2002, "Application of Shared Telerobotic Control for Contour Following Process," *Australasian Conference on Robotics and Automation*, November, pp. 88–93.
[31] Griffin, W. B., 2003, *Shared Control for Dexterous Telemanipulation with Haptic Feedback*, Stanford University, Stanford, CA.
[32] Griffin, W. B., Provancher, W. R., and Cutkosky, M. R., 2003, "Feedback Strategies for Shared Control in Dexterous Telemanipulation," *Intelligent Robots and Systems, IROS 2003, Proc. 2003 IEEE/RSJ Int. Conf.*, IEEE, New York, pp. 2791–2796.
[33] Payandeh, S., 2001, "Application of Shared Control Strategy in the Design of a Robotic Device," *Proceedings of the 2001 American Control Conference*, June 25–27, IEEE, Arlington, VA, pp. 4532–4536.
[34] O'Malley, M. K., and Gupta, A., 2003, "Passive and Active Assistance for Human Performance of a Simulated Underactuated Dynamic Task," *International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE, Los Angeles, CA, March, pp. 348–355.
[35] O'Malley, M. K., and Gupta, A., 2003, "Skill Transfer in a Simulated Underactuated Dynamic Task," *International Workshop on Robot and Human Interactive Communication*, IEEE, Millbrae, CA, Oct.–Nov., pp. 315–320.
[36] Jagacinski, R., and Flach, J., 2003, *Control Theory for Humans: Quantitative Approaches to Modeling Performance*, Lawrence Erlbaum Associates, Mahwah, NJ, p. 379.