v 2.0, 5 May 2004
Submitted to IEEE Trans. Neural Syst. Rehabil.

Human adaptation to dynamic interaction forces during learning of a visuo-motor task

Felix C. Huang, R. Brent Gillespie, and Arthur D. Kuo
Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109 USA
e-mail: fhuang@umich.edu

Key words: haptic feedback, skill transfer, virtual reality, motor control, upper extremity, feedback control

Abstract

We tested whether humans can learn to sense and compensate for interaction forces in contact tasks. Many tasks, such as use of hand tools, involve significant interaction forces between hand and environment. One control strategy would be to use high hand impedance to reduce sensitivity to these forces. But an alternative would be to learn feedback compensation for the extrinsic dynamics and associated interaction forces, with the potential for lower control effort. We observed subjects as they learned control of a ball-and-beam system, a visuo-motor task where the goal was to quickly position a ball rolling atop a rotating beam, through manual rotation of the beam alone. We devised a ball-and-beam apparatus that could be operated in a real mode, where a physical ball was present; or in a virtual training mode, where the ball's dynamics were simulated in real time. The apparatus presented the same visual feedback in all cases, and optionally produced haptic feedback of the interaction forces associated with the ball's motion. Two healthy adult subject groups, Vision-Only and Vision-Haptics (each N=10), both trained for 80 trials on the simulated system, and then were evaluated on the real system to test for skill transfer effects. If humans incorporate interaction forces in their learning, the Vision-Haptics group would be expected to exhibit a smoother transfer, as quantified by changes in completion time of a ball-positioning task. During training, both groups adapted well to the task, with reductions of 64-70% in completion time. At skill transfer to the real system, the Vision-Only group had a significant 35% increase in completion time ($p < 0.05$). There was no significant change in the Vision-Haptics group, indicating that subjects had learned to compensate for interaction forces. These forces could potentially be incorporated in virtual environments to assist with motor training or rehabilitation.

1. Introduction

Sensory information is important for both learning and execution of many motor tasks. It is critical for tasks that require continuous adjustment of position or maintenance of balance, such as standing upright (Park et al. 2003) or holding a cup of coffee (Dingwell et al. 2002). Some actions, such as in touch-typing, occur too quickly for feedback to contribute directly to execution (Rempel et al. 1994), but sensory information is nonetheless important for assessing outcome and providing re-afference for motor learning (Johansson and Cole 1992). It might therefore be helpful to provide sensory feedback to humans learning novel motor tasks or undergoing neuro-rehabilitation (Patton and Mussa-Ivaldi 2003). Virtual environments can produce a wide range of visual and haptic (touch) feedback, with the potential to manipulate that feedback and perhaps enhance learning. But before this potential can be realized, it is important to understand how the quality and quantity of sensory feedback contributes to learning and performance.

In many motor tasks, interaction forces are important for feedback control. For reaching tasks, the main feedback is proprioceptive and visual. Muscle spindles, Golgi tendon organs, and nociceptors provide proprioceptive feedback of muscle length, muscle force, and joint position, respectively (Gandevia et al. 2002). Vision is often important for continuous feedback and knowledge of results, particularly for visually-guided tasks where the hand must be positioned relative to objects in the environment (Elliott et al. 1998). But many skilled tasks also involve contact and interaction with the environment, for example, washing a window or turning a crank. In such tasks, haptic feedback from mechanoceptors allows the interaction forces with the environment to be monitored (Westling and Johansson 1987) and potentially regulated.

The central nervous system (CNS) integrates this array of information and learns to associate sensory patterns with motor actions. Some investigators have hypothesized that this sensory integration is part of a CNS internal representation of movement (e.g., Cosman et al. 2002), where associations are developed between a motor command and the resulting motion. This internal representation may take the form of a *forward model* of limb dynamics, where the motor command can be used to predict the expected motion and associated sensory feedback (Johansson and Cole 1992; Mehta and Schaal 2002). Differences between expected and actual afference might then be used to refine the motor command during learning. Another form of internal representation is an *inverse model* of limb dynamics (Kawato 1999), where a desired outcome is used to trigger the appropriate motor command for the task. Both types of internal representation are learned with the help of sensory feedback (Wolpert and Kawato 1998), and model-based control allows movements to be more automatic and to rely less heavily on high-gain feedback (Shadmehr and Mussa-Ivaldi 1994). This learning can compensate not only for the natural limb dynamics, but also altered dynamics including viscous and spring-like environments applied by robotic manipulators (Scheidt et al. 2001; Dingwell et al. 2002). These manipulators alter the torques necessary to perform a reaching command, and subjects' adaptations demonstrate that humans can adapt to a wide variety of environments.

A CNS internal representation could potentially include not only dynamics of the limbs, but of the external environment as well. Humans appear to have some predictive visuo-motor understanding of the kinematics of external objects, as is needed to catch a ball (McIntyre et al. 1998) or balance a pole on one's hand (Mehta and Schaal 2002). This requires sensing and prediction of kinematical states external to the body. Humans also appear to haptically and proprioceptively sense the mass of an external object simply by wielding it (Chan 1994; Cosman et al. 2002). Depending on how the object is coupled to the hand, it may simultaneously introduce dynamic interaction forces in addition to extrinsic kinematic states. Compensation for such forces was demonstrated by Dingwell et al. (2002), in a fast reaching task where subjects exerted forces through a virtual spring to position a virtual mass on the opposite end. After a training period, subjects adapted well to the spring-mass dynamics, producing motions that could not be explained by increased hand impedance or greater use of visual feedback. The motions were best explained by an internal representation of the spring-mass dynamics, allowing subjects to produce an appropriate feedforward control. This leads to the question of whether interaction forces are also incorporated in learning of a feedback control task, in which extrinsic states must be continuously stabilized. Such learning may be relevant to motor training or rehabilitation, because interaction forces are important in many skilled occupational tasks involving tool use.

The purpose of the present study was to test whether humans learn to accommodate or counteract interaction forces in a novel motor task that requires feedback control of states external to the body. We performed this test with the *ball-and-beam task* (Fig. 1A), where the goal is to position and balance a ball rolling freely atop a beam that pivots about a horizontal axis, by manually rotating the beam. Because the ball's position and velocity are not directly governed by the hand, these constitute states external to the body. As with pole-balancing (Mehta and Schaal 2002), the ball-and-beam system is dynamically unstable. While the ball's motion can be initiated with a feedforward beam rotation, feedback is necessary to maintain a target ball position. The ball's motion can be sensed visually, and to a greater degree than pole balancing, also haptically. This task could in principle be performed by learning the kinematics of the ball and beam, and using hand positioning with high impedance to rotate the beam appropriately. But incorporation of a control strategy specific to the ball-and-beam dynamics could allow for lower hand impedance and minimal effort. We tested for learning of interaction forces with a skill transfer experiment, where subjects learned to control the ball in a virtual environment, with sensory feedback as an independent variable. After training one subject group with visual feedback alone and another with visual and haptic feedback, we assessed their task performance on a physical ball-and-beam system. We hypothesized that subjects trained with haptic feedback would learn the interaction forces and would subsequently exhibit better skill transfer to the physical task compared to subjects receiving visual feedback only.

2. Methods

We developed a ball-and-beam motor task trainer and simulator (Fig. 1A), with which subjects learned how to balance a ball atop a pivoting beam. The apparatus could be operated in two modes, where the subject held the beam at one end and manually rotated it to influence the motion of either 1) a physical ball rolling on the beam, or 2) a virtual ball whose rolling was simulated by computer. In the first case, the subject could directly feel the interaction of the physical ball with the beam. In the second, the virtual ball-and-beam dynamics were simulated in real-time, and the interaction forces presented via the beam to the subject's hand by a motor driving the beam. In both modes, visual feedback was provided by a computer graphic animation of the ball and beam. We trained two groups of subjects to move the virtual ball to a target position on the beam, wherein one group was provided with visual feedback alone, and the other group was provided with visual and haptic feedback. After a training period, the performance of both groups was evaluated using the physical ball atop the same beam. We then tested for adaptation to interaction forces, by observing how well each group transferred to the physical system. The following sections describe the dynamics of the ball-and-beam system, its implementation in a virtual environment, and the protocol of the skill transfer experiments.

2.1 Virtual Ball-and-Beam Environment

We modeled the ball-and-beam system (Fig. 1B) with two coupled second-order equations of motion (Eqs. 1-2). A beam with transverse moment of inertia I is connected to a base, pivoted about its midpoint. A ball of mass m and radius R rolls without slipping over the center-line of the beam. The angular displacement of the beam θ is measured counter-clockwise relative to the ground. The translational displacement of the ball relative to the beam center is r, increasing to the right. The ball's mass is subject to gravity g, and its moment of inertia is $\alpha m R^2$, where α = 2/5 for a sphere. Two torques act on the ball and beam system: the motor torque τ, and the torque due to the interaction force F acting between the subject's hand and the

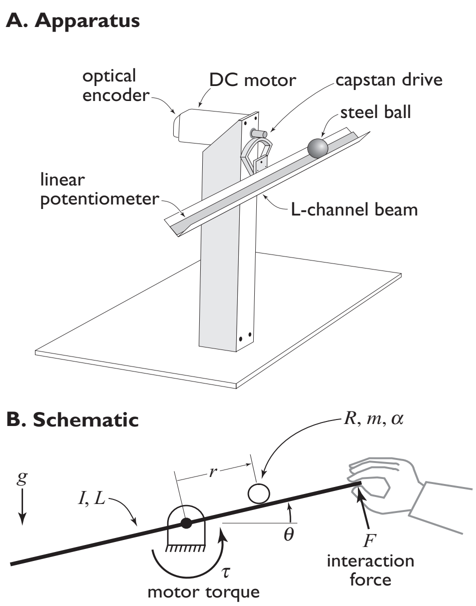

Figure 1. (A.) In the ball-and-beam task, a human operator rotates a beam, atop which a ball rolls without slipping. The goal of the task is to control the ball by moving the beam alone. The ball's motion is determined by the tilt of the beam. The human operator experiences interaction forces at the point of contact, with the forces depending on the motion of the ball. The ball-and-beam apparatus may be operated with a physical ball, or with a virtual representation of the ball, with feedback of interaction supplied by a DC-motor. (B.) Schematic shows definition of mathematical symbols (Table 1). The subject exerts interaction force F against the end of the beam, and the computer-driven motor exerts torque τ.

beam. We obtained the equations of motion for this system using Lagrangian methods:

$$0 = (1+\alpha)\ddot{r} - R\ddot{\theta} - r\dot{\theta}^2 + g\sin\theta \qquad \textbf{Eq. 1}$$

$$F \cdot \frac{L}{2} + \tau = -m\alpha R\ddot{r} + \left[I + m\left(\alpha R^2 + r^2\right)\right]\ddot{\theta} + 2mr\dot{r}\dot{\theta} - mg\left(r\cos\theta + R\sin\theta\right) \qquad \textbf{Eq. 2}$$

The equations above comprise the model used to simulate the ball-and-beam dynamics with real-time numerical integration. The motion of the ball along the beam is determined solely by the angular input $\theta(t)$ and its derivatives (Eq. 1), whereas the interaction force F depends on motion of both ball and beam (Eq. 2). Linearized about a non-zero equilibrium position r, the system has two unstable poles. To maintain such a position, it is necessary to produce compensatory beam motions dependent on feedback of $r(t)$, $\theta(t)$, and their time-derivatives. When the system is operated with the physical ball, the motor is turned off, and so τ = 0. When operated with the virtual ball, the motor is programmed to supply the torques associated with the missing physical ball (all terms containing the ball mass m on the right-hand-side of Eq. 2).

2.2 Description of Apparatus

Our experimental apparatus consists of a lightweight beam, pivoted about a horizontal axis, atop which a ball could roll without slipping (see Fig. 1A). The beam was constructed from L-channel aluminum, mounted with the inner faces directed upward and connected to a support tower via a revolute joint with horizontal axis. For the ball, we used a chrome steel ball bearing rolling within the L-channel. The beam's axis of rotation was designed to pass through the path of the ball's mass center. Actuation of the beam by the motor occurred through a capstan drive to an aluminum arc fixed to the beam via an inextensible flexible steel cable. The capstan drive supplied a mechanical advantage of 11.68 between the motor and the beam. The relevant inertial and geometric parameters of the apparatus are shown in Table 1.

Table 1. Inertial and geometric properties for the ball-and-beam apparatus.

Beam Moment of Inertia	I	0.027 kg-m^2
Beam Length	L	0.80 m
Ball Mass	m	0.23 kg
Ball Radius	R	0.019 m

The experimental apparatus was used to collect data during the ball-and-beam task, and in some conditions to display interaction forces in the form of haptic feedback without the physical ball. In the haptic feedback conditions, the DC motor (peak stall torque 0.30 N-m) replaced the torques normally felt by the operator when the physical ball was present. A linear potentiometer transduced the position of the steel ball along the beam. The potentiometer consisted of a plastic conductive strip on an inner face of the beam, acting as a variable resistance in a voltage divider circuit, with the ball making contact between the strip and the other face of the beam. The measured voltage was linearly proportional to the ball position, with a resolution of approximately 0.0005 m. The beam's angle of rotation was measured with a quadrature-output optical encoder (3600 lines per revolution) mounted on the motor shaft. A desktop PC collected these data and controlled the motor in real-time, with a sampling rate of 1 kHz. Experimental data were logged at a lower rate of 100 Hz.

2.3 Implementation of Virtual and Real Environments

The ball-and-beam apparatus could function in *real mode* using a physical ball, and also in *virtual mode* using a virtual ball with simulated dynamics (see Fig. 2). In both modes, the human operator moved a physical beam and received visual feedback from a real-time animation of the ball-and-beam on a computer monitor. In real mode, the human operator actuated the beam with the motor turned off, with the position of the physical ball and orientation of the angle recorded by computer and displayed in the animation. In virtual mode, the physical ball was not present, and in its place the computer simulated the dynamics and optionally provided haptic feedback in the form of interaction torques associated with the ball's motion.

In virtual mode, the computer measured the human operator's input to the ball-and-beam system, in the form of the beam angle θ. The ball motion was simulated in real time, along with the interaction force that would have been produced by the physical ball had it been present. The simulation was performed by integrating a discretized form of Eq. 1 in real time. This yielded the ball motion which was both displayed in the animation and logged as data. In addition to this visual feedback, the computer optionally provided haptic feedback, by computing the interaction torque from Eq. 2, and supplying this torque through the DC motor.

The programmed dynamics were designed to emulate the force coupling characteristics of the real rolling ball, which we consider to be the most critical element of haptic feedback in the real system. However, for this experiment, we did not include other effects which may be additional sources of feedback to a human operator. The effects of rolling irregularities and friction may produce both haptic and audio feedback, which were not included in our virtual environment.

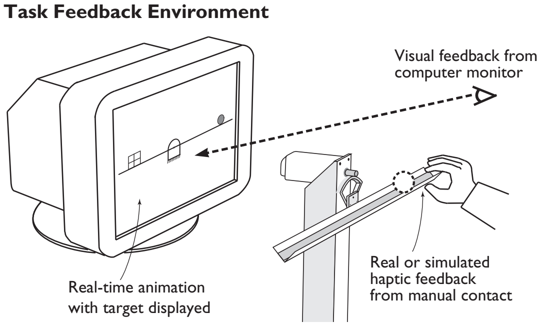

Figure 2. Each human subject received visual feedback from the computer animation including the display of a box for the target ball position. Simulated haptic feedback from a motor or from the presence of a real ball produced interaction forces between the beam and hand.

2.4 Transfer Paradigm Experiment Protocol

Our experiment compared skill transfer in a ball balancing task for two subject groups, Vision-Only and Vision-Haptics, as outlined in Table 2. Twenty young healthy adults (ages 22 – 39) were randomly assigned to the two subject groups. All participants provided informed consent in accordance with University of Michigan human subject protection policies.

Table 2. Haptic feedback conditions applied during training and evaluation sessions with two subject groups. The Vision-Only group learned the ball-and-beam task (Training Session) without haptic feedback, and then performed the task (Evaluation Session) with the physical ball. The Vision-Haptics group received virtual haptic feedback during Training, before performing the task with the physical ball. Visual feedback was animated by computer in all trials.

Subject Group	Training Session w/ Virtual Ball (80 Trials)	Evaluation Session w/ Real Ball (80 Trials)
Vision-Only (n=10)	No haptic feedback	Physical haptic Feedback
Vision-Haptics (n=10)	Virtual haptic feedback	Physical haptic Feedback

Both groups trained on a virtual system, after which their performance was evaluated on the physical system. Subjects from both groups viewed an animation of the ball-and-beam motion on a computer screen as they performed the ball balancing task. All training and evaluation trials presented visual feedback. The

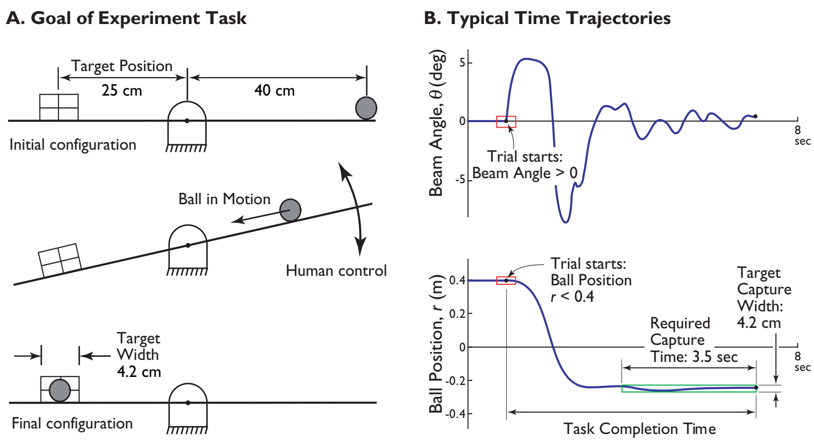

Figure 3. (A.) The goal of the experiment task was to move the beam so that the ball would move from the far right hand end of the beam to a target box in a minimum amount of time. The *task completion time* was defined as the time from the beginning of ball motion to the instant at which the ball is settled within the target zone for a required capture time of 3.5 sec. (B.) Typical trajectories of beam angle and ball position, from one subject.

two groups differed as to whether they received haptic feedback during training. Vision-Haptics subjects received haptic feedback supplied by the computer, to simulate the interaction forces associated with ball motion. The Vision-Only did not receive any haptic feedback of interaction forces.

We employed a skill transfer paradigm with a training session of 80 practice trials, and then an evaluation session of another 80 trials. Training took place with a virtual ball, and evaluation with the physical ball. We asked subjects to complete a ball positioning movement in minimum time for all trials. A short break (≈1 min.) was given after every 20 trials to avoid fatigue. Subjects did not receive an extended rest period between training and evaluation sessions. Subjects typically completed both sessions in approximately 35 min.

The goal of the task was to control the beam so that the rolling ball would quickly move from one end of the beam (r = 40 cm) to a target box (4.2 cm wide) across the beam (r = -25 cm), and then keep the ball center within the target box for a minimum capture time of 3.5 sec (see Fig. 3). Positive displacement of the ball was defined to the operator's right, when facing the beam. The target box was displayed along with the animated ball-and-beam. The ball was initially at rest, with each trial commencing when the subject moved the beam counter-clockwise from the horizontal, thereby causing the ball to start rolling. Audio beeps and a color cue displayed on the monitor alerted subjects when the ball was settled within the target box for the required capture time. The *task completion time* was defined from onset of ball motion to completion of this capture. After completion, subjects were instructed to roll the ball back to the starting position, at which point the computer would indicate readiness for the next trial. Subjects were asked to achieve the lowest task completion time possible for each trial.

Each participant was given the same instructions for body stance and hand positioning on the beam and the criteria for performance was explained. Body stance consisted of a standing position with the left hand resting on the table surface. Each subject controlled the motion of the beam using the right hand, holding onto the end of the beam. We used the same hand for all subjects, regardless of handedness (1 out of 20 subjects were self-reported as left-handed) because the task was fairly novel and we were interested mainly in each subject's improvement in control regardless of initial level of performance, For all trials, subjects viewed an animation of the ball-and-beam motion on a computer monitor without looking at the physical apparatus. There was no prompting of strategy to the participants, nor was there a demonstration of the control action.

2.5 Task Completion Time and Analysis of Skill Transfer

In order to determine differences in performance between the two subject groups, we undertook two types of data analysis. The first was a simple comparison of absolute task completion times, to form a non-parametric indicator of learning during the training and/or evaluation sessions. This comparison was based on block-averaged task completion times, formed by averaging 20 trial blocks at the start and end of each session. We performed an analysis of variance on these block-average task completion times, considering three experiment factors: training group (Vision-Only vs. Vision-Haptics), testing session (Training vs. Evaluation), and testing period (blocks from Start or End of the testing session). We also considered two-way interactions, with significant interactions leading to a limited set of post-hoc tests to determine general indications of learning. We used paired, two-tailed t-tests with $p < 0.05$ as the criterion for statistical significance.

The second analysis used a parametric fit to normalize and characterize the learning curves, and determine whether there were relative differences in skill transfer. The parametric fit was a power law, as a function of trial number (see Fig. 4). This fit was performed on each subject's task completion time data, with separate fits for the training and evaluation sessions. These fits were then used to compare the change in performance during skill transfer, by examining the *transition delta*, or change in completion times when switching from a virtual system to the physical system. Comparisons of transition deltas between groups were then made using paired, two-tailed t-tests using $p < 0.05$ as the criterion for statistical significance.

To account for differences between individuals when performing the parametric fit, task completion times were scaled by a normalizing factor specific to each subject. This normalizing factor, the *final-performance norm*, was defined as the average task completion time of the last twenty trials from each subject's evaluation session (see Fig. 4). Normalized task completion time data were determined by dividing all of the subject's task completion times by the final performance norm.

The normalized task completion time data from the training and evaluation sessions were then fit, separately, to power law functions using a non-linear least squares method (Trust-Region algorithm from MATLAB Curve-Fitting Toolbox, MathWorks, Natick, MA), where the parameters were found using a

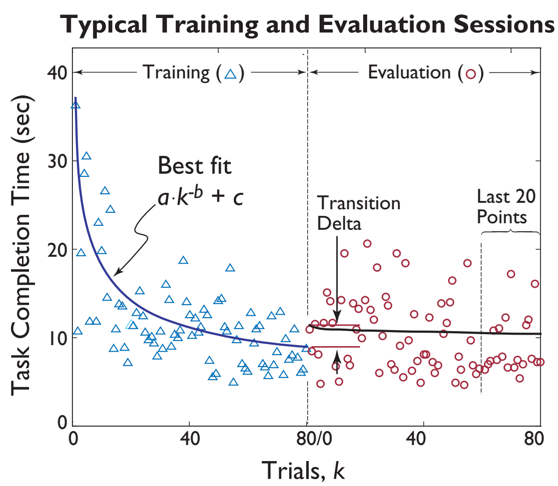

Figure 4. Results from a typical subject's training and evaluation sessions. Task completion times are plotted as a function of trial number. After 80 trials of training on the virtual ball-and-beam, subjects were evaluated on the real ball-and-beam. Power law curve fits describe the overall learning trends. The transition delta quantifies the skill transfer in terms of the difference between curves at transition from training to evaluation. The last 20 points of the evaluation session were used as a normalization factor for subsequent analysis.

bounded search. The power law functions used, typical of skill acquisition (Newell and Rosenbloom 1981), were of the form

$$T_{\text{task}}(k) = a \cdot k^{-b} + c\,, \qquad \textbf{Eq. 3}$$

where $T_{\text{task}}(k)$ was the task completion time for the k-th trial, starting from $k = 1$. Parameter a weights the function to have proper initial value, b indicates the rate of convergence of the curve, and c is an offset representing the limit of task completion time as k increases.

These performance curves represent the overall trends of task completion time as the number of trials progressed. The curves demonstrate a general reduction in task completion time with practice. Sample task completion results, shown for one subject in Figure 4, illustrate two separate best-fit curves overlaid onto the normalized completion time data. The first curve starts at a large value and decreases as the training trials progress. The second curve starts at a value higher than the end of the previous curve, but approaches the final performance norm as the evaluation trials progress.

We examined the success in skill transfer for each subject by considering the *transition delta*, defined as the fitted completion time (Eq. 3) evaluated at the beginning of the evaluation session, minus that from the end of the training session (Figure 4). A positive transition delta indicated an initial deficit in performance in the evaluation session. We compared the fitted performance curves and transition delta results from the Vision-Only and Vision-Haptics groups, to determine what differences arose from the groups' respective training conditions. As with block averages, these comparisons were based on t-tests with $p < 0.05$ as the criterion for significance.

3.0 Results

Subjects in both groups improved their performance of the ball positioning task with practice. The progression of task completion times through trials was characterized by a steep initial descent and subsequent leveling, as shown in the sample results of Figure 4. A reduction in the scatter of task completion times was typical of both subject groups as well. Most subjects demonstrated convergence of performance mainly within the first 60 to 80 training trials, but with continuing small improvements through the end of the evaluation session, through the total 160 trials.

Block averaged task completion times revealed significant differences associated with each factor. Analysis of variance yielded significant differences for subject training group (Vision-Only vs. Vision-Haptics, p = 0.0068), testing session (Training vs. Evaluation, p = 1.2e-7, and trial block (Start vs. End, p = 3.5e-6). Two-way interactions were significant for testing session and trial block (p = 9.8e-6), and for subject training group and trial block (p = 0.012).

Both subject groups demonstrated a marked decrease in the block-averaged task completion times during the training session (Table 3). Quantitative improvements in performance with training were found to be statistically significant, with reductions of at least 30-40% (Vision-Only group, p=0.00024; Vision-Haptics group, p=0.00036). The group standard deviations also decreased from the start to the end of the training session for both groups, indicating more consistent trial-to-trial performance with practice.

Although the major adaptations occurred during training, some subjects also exhibited significant changes in performance during the evaluation session. The Vision-Only group had a significant, 2 second decrease in block averaged task completion times (p=0.020) with continued practice, whereas the Vision-Haptics group demonstrated no significant change (p=0.14) through the evaluation session. The practice gained during the evaluation session was sufficient to result in no significant differences in performance between the two groups (p=0.90) by the end of the evaluation session. The average completion times at the end of the evaluation session (last column of Table 3) also serve as the final performance norms for use in the power law fits. The experiment-wise error rate for these post-hoc test results was p = 0.021.

Table 3. Block-averaged absolute task completion times (in sec) for the start and end of training and evaluation sessions. Start and End values shown are means (± s.d.) taken over the first or last 20 trials, respectively, of each session. Significant differences were observed for each of three factors (p < 0.05).

Group	Training Session		Evaluation Session	
	Start	End	Start	End
Vision-Only	21.0 ± 7.8	12.4 ± 3.7	12.3 ± 3.1	10.3 ± 2.5
Vision-Haptics	17.6 ± 4.8	12.0 ± 2.6	10.7 ± 1.8	10.1 ± 1.8

Power law curves, fit to each subject's normalized task completion times, appeared to successfully capture the trends of the raw task completion data points. Mean learning curves (Fig. 5) demonstrated decreases of 64-70% in task completion time, as well as decreased variability (group standard deviations) with training. The Vision-Haptics group began training with insignificantly different task completion times compared to Vision-Only, 3.16±1.66 versus 4.04±1.80 (p=0.27), and also finished the training session with comparable performance, 1.13±0.11 versus 1.19±0.18 (p=0.42), in units of normalized time. Upon switching to the physically real system, the Vision-Haptics group began the evaluation session with a significantly lower completion time 1.19±0.23 versus 1.61±0.43 (p=0.016). In agreement with our analysis of the final performance norms, the subject groups finished evaluation with comparable performance (0.95±0.05 vs. 0.97±0.07, no significant difference p=0.87). In each case, standard deviations were smaller for the Vision-Haptics group through the latter part of training and throughout the evaluation session. The power law parameters, from which these results are derived, are shown in Table 4.

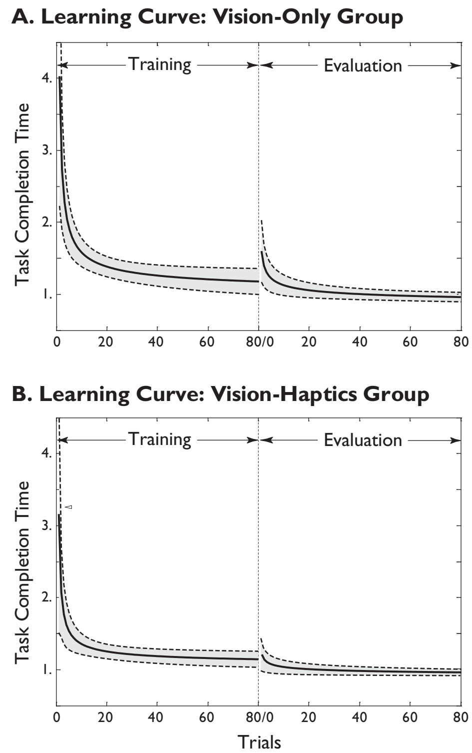

Figure 5. Learning curves for (A.) Vision-Only and (B.) Vision-Haptics groups, during Training and Evaluation Sessions. Learning curves shown are derived from power law curves (Eq. 3), averaged over subjects in each group (region within dashed lines denote ±1 standard deviation). Both groups adapted to the virtual ball-and-beam task during the Training Session, and then experienced a small increase in task completion time when transferring to the Evaluation Session with the physical ball. Further adaptation was observed as the Evaluation Session progressed. Vertical axis is task completion time, normalized by the final performance time.

Table 4. Power Law Curve Fit Parameters (Eq. 3) and R-Squared Values (mean ± s.d.).

Group	Parameter	Training Session	Evaluation Session
Vision-Only	a	3.17 ± 1.54	1.09 ± 0.51
	b	0.75 ± 0.42	0.35 ± 0.30
	c	0.87 ± 0.55	0.52 ± 0.44
	R^2	0.59 ± 0.11	0.38 ± 0.19
Vision-Haptics	a	2.42 ± 1.28	0.67 ± 0.50
	b	0.79 ± 0.71	0.28 ± 0.30
	c	0.74 ± 0.53	0.53 ± 0.46
	R^2	0.61 ± 0.14	0.34 ± 0.18

In these performance curves, the Vision-Haptics group exhibited a better skill transfer than the Vision-Only group at the point of transition from the training to the evaluation session (Fig. 5). The performance curves of the Vision-Haptics group were fairly continuous at the point of transition. For the Vision-Only group, however, a large positive offset, or transition delta, was evident between the performance curves of the two sessions. These results are expressed quantitatively by a comparison of normalized task completion times derived from power law curve fits (Fig. 6A). The Vision-Haptics group demonstrated an insignificant change in normalized task completion time (1.13±0.11 to 1.20±0.23, p=0.36), when switching to the real ball-and-beam. In contrast, the Vision-Only group demonstrated a significant increase in task completion time after transition (p=0.0028), from 1.19±0.18 to 1.61±0.43, indicating a disruption in performance. A comparison of the changes in performance (see Figure 6B) shows that the Vision-Haptics group exhibited a significantly smaller change (p=0.011), 0.07±0.22 vs. the Vision-Only group's 0.42±0.33.

4.0 Discussion

We sought to test whether subjects compensated for interaction forces between the hand and the beam. The ability to learn control of extrinsic states has previously been shown in visuo-motor tasks, both in feedforward and feedback settings. In a feedforward task, adaptation to dynamics is readily demonstrated with "catch trials," where the external object is unexpectedly altered, and resulting changes in ballistic movement belie compensation for the original dynamics (Dingwell et al. 2002). In a feedback task, the adaptation can be probed by altering the dynamics of the object or the sensory information regarding the object's motion. For example, temporal features of feedback compensation can be characterized by blanking out sensory information and observing the time course of the motor command (Mehta and Schaal 2002). In the present study, the experimental alteration consisted of switching from one of two virtual environments to the physical environment and then observing changes in performance, both immediately and over time.

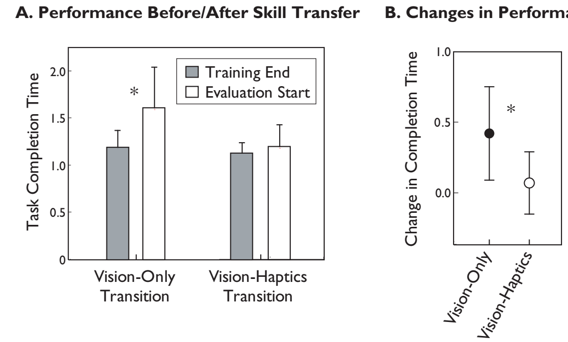

Figure 6. (A.) Performance and (B.) changes in performance at skill transfer, comparing Vision-Only and Vision-Haptics groups. Data are derived from learning curves (Fig. 5), showing mean task completion time (error bars denote one s.d.). At skill transfer, Vision-Only group demonstrated significantly degraded performance in terms of longer completion time compared to end of Training Session. Vision-Haptic group had no significant change in performance. Comparing the two groups directly (B.), the increase in completion time at skill transfer was significantly larger for Vision-Only vs. Vision-Haptics. For both graphs, vertical axis is time, normalized by the final performance time. (* indicates $p < 0.05$.)

Our results show evidence that subjects do incorporate interaction forces when learning to control the ball and beam. This is demonstrated by the disparate results for the two subjects groups when evaluated for skill transfer (Fig. 6). The Vision-Haptics group, which had been exposed to simulated interaction forces, achieved a relatively smooth skill transfer, with only an insignificant 6.2% increase in completion time on the physical system. In contrast, Vision-Only subjects learned to control visual aspects of the task equally well, but experienced a significant 35% increase in completion time when transferring to the real system. Even though both groups were able to control the ball-and-beam system equally well by the end of the training session, the Vision-Haptics group was better able to compensate for the interaction forces felt in the real system during the evaluation session.

These results are in contrast to what would be expected if subjects had only learned to control the kinematics of the system, that is, to associate desired ball motion with a beam position command, but not with the interaction forces necessary to achieve that beam position. Kinematic control could be achieved by controlling the beam with high hand impedance, but such control would be negligibly affected by interaction forces, and would be expected to produce a seamless skill transfer. The ability to control the ball certainly indicates an acquired understanding of the ball-and-beam kinematics. Indeed, Vision-Only subjects performed far better at the onset of skill transfer than they had when exposed to the task for the first time, suggesting that kinematic training alone was helpful for learning. But the more effective skill transfer demonstrated by subjects trained with haptic feedback indicates the successful incorporation of interaction forces in their control strategy.

The non-zero change in performance seen in the Vision-Haptics group indicates that the virtual ball-and-beam system, despite its advantages over Vision-Only, was nevertheless an imperfect model of reality. A sufficiently realistic model would be expected to yield an insignificant change in performance at skill transfer. There were a number of aspects of the real system not replicated in the virtual system, such as the fact the that real-time ball-and-beam model is limited in bandwidth and can only approximate the actual interaction forces. But more important differences may lie in other sensory dimensions missing from the virtual system. Some subjects noted that they found the sound of the physical ball rolling on the beam to be helpful. This rolling also produced a very light but detectable vibration in the beam, that may provide additional cues regarding ball velocity. With regard to the experiment, low fidelity of simulation would be expected to reduce the magnitude of differences observed between groups. At any rate, the skill transfer may be a useful paradigm for evaluating the fidelity of a virtual reality system.

Interaction forces are useful for control, partly due to the sensory information they impart. The Central Nervous System (CNS) has a large array of sensory information available, and it is advantageous to use any sensory cues that might contribute to estimation of the state of the body and environment. In reaching tasks, visual and proprioceptive feedback is sufficient to estimate limb position and speed. Contact tasks may involve the same limb state but vastly different interaction forces, and haptic information can help to differentiate between such forces. In the ball and beam task, interaction forces are nearly proportional to ball position, and can therefore serve almost directly as a measurement of extrinsic states.

Unlike other sensory feedback, however, interaction forces can disturb the hand. This is due to the two-way dynamic coupling between hand and beam, which involves both position and force. In the presence of these interaction forces, a desired beam position can only be achieved if the CNS uses high impedance control to make the hand stiff, or produces appropriate compensatory forces. High impedance reduces sensitivity to disturbances, but is also potentially costly in terms of muscle force or activation, and might require co-contraction. A gradual reduction in impedance and/or co-contraction has been observed in a number of other motor tasks (Thoroughman and Shadmehr 1999; Osu et al. 2002; Gribble et al. 2003). In contrast, learning of compensatory forces reduces the need for co-contraction, and instead relies on both intrinsic and extrinsic states, whose dynamics are likely learned over time. It is possible that subjects initially used higher impedance at the beginning of the Training Session, but gradually reduced that strategy over time. If so, haptic feedback may be particularly advantageous for learning compensatory control.

Our results are not, however, sufficient to disprove the use of high-impedance strategies. Here, it is possible that Vision-Haptics subjects used haptic feedback simply to select higher impedance rather than to learn the ball-and-beam interaction forces, with vision alone sufficient to learn ball-and-beam kinematics. But such impedance tuning would be expected to occur very quicly after skill transfer, in contrast to the continued adaptation we observed throughout the evaluation sessions.

The short-term advantages of learning with haptic feedback do not, however, appear to be critical in the long term. By the end of the Evaluation Session, the two subject groups performed the task nearly equally well, with no significant difference in performance. The presence or absence of haptic feedback therefore does not appear to interfere with learning during the Training Session, nor with the eventual steady-state performance at the end of the Evaluation Session. The main advantage of haptic feedback was to effect smoother skill transfer than with vision alone.

These findings may have implications for rehabilitation applications. Virtual environments are potentially useful for rehabilitation because they can apply highly customized and repeatable training, while simultaneously quantifying performance as subjects improve. Because many skilled occupational tasks involve contact with and manipulation of the environment, it is potentially advantageous for virtual environments to simulate these tasks, and to provide haptic feedback to help subjects to learn the associated interaction forces.

Acknowledgements

This work supported in part by the Midwest Regional Rehabilitation Network, funded by NIH 1R24 HD39627-01. Peter Lin provided invaluable engineering assistance with the ball and beam virtual environment.

References

Chan TC (1994) Haptic perception of partial-rod lengths with the rod held stationary or wielded. Percept Psychophys 55: 551-561

Cosman PH, Cregan PC, Martin CJ, Cartmill JA (2002) Virtual reality simulators: current status in acquisition and assessment of surgical skills. ANZ J Surg 72: 30-34

Dingwell JB, Mah CD, Mussa-Ivaldi FA (2002) Manipulating objects with internal degrees of freedom: evidence for model-based control. J Neurophysiol 88: 222-235

Elliott D, Ricker KL, Lyons J (1998) The control of sequential goal-directed movement: learning to use feedback or central planning? Motor Control 2: 61-80

Gandevia SC, Refshauge KM, Collins DF (2002) Proprioception: peripheral inputs and perceptual interactions. Adv Exp Med Biol 508: 61-68

Gribble PL, Mullin LI, Cothros N, Mattar A (2003) Role of cocontraction in arm movement accuracy. J Neurophysiol 89: 2396-2405

Johansson RS, Cole KJ (1992) Sensory-motor coordination during grasping and manipulative actions. Curr Opin Neurobiol 2: 815-823

Kawato M (1999) Internal models for motor control and trajectory planning. Curr Opin Neurobiol 9: 718-727

McIntyre J, Berthoz A, Lacquaniti F (1998) Reference frames and internal models for visuo-manual coordination: what can we learn from microgravity experiments? Brain Res Brain Res Rev 28: 143-154

Mehta B, Schaal S (2002) Forward models in visuomotor control. J Neurophysiol 88: 942-953

Newell A, Rosenbloom PS (1981) Mechanisms of skill acuisition and the power law of practice. In: Anderson JR (ed) Cognitive skills and their acquisitions. Erlbaum, Hillsdale, NJ

Osu R, Franklin DW, Kato H, Gomi H, Domen K, Yoshioka T, Kawato M (2002) Short- and long-term changes in joint co-contraction associated with motor learning as revealed from surface EMG. J Neurophysiol 88: 991-1004

Park S, Horak FB, Kuo AD (2003) Postural feedback responses scale with biomechanical constraints in human standing. Exp Brain Res: in press

Patton JL, Mussa-Ivaldi FA (2003) Robot-assisted adaptive training: A new paradigm for motor learning. IEEE Transactions on Biomedical Engineering: in press

Rempel D, Dennerlein J, Mote CD, Jr., Armstrong T (1994) A method of measuring fingertip loading during keyboard use. J Biomech 27: 1101-1104

Scheidt RA, Dingwell JB, Mussa-Ivaldi FA (2001) Learning to move amid uncertainty. J Neurophysiol 86: 971-985

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14: 3208-3224

Thoroughman KA, Shadmehr R (1999) Electromyographic correlates of learning an internal model of reaching movements. J Neurosci 19: 8573-8588

Westling G, Johansson RS (1987) Responses in glabrous skin mechanoreceptors during precision grip in humans. Exp Brain Res 66: 128-140

Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. Neural Netw 11: 1317-1329