Progressive Shared Control for Training in Virtual Environments

Yanfang Li * Rice University
Joel C. Huegel † Rice University
Volkan Patoglu‡ Sabancı University, Istanbul, Turkey
Marcia K. O'Malley§ Rice University

ABSTRACT

Virtual environments (VEs) with haptic feedback not only provide a safe and versatile practice medium for many manual control tasks, but also promise to improve the efficacy of manual skill training with the addition of haptic guidance. However, haptic guidance schemes such as shared control may be detrimental since such schemes actively interfere with the coupled system dynamics, thereby causing participants to experience task dynamics that are altered from those of the real task. Therefore, this paper proposes a *progressive approach* to guidance for training in virtual environments. This progressive guidance scheme adjusts its control gains based on participant performance, aiming to expose the performer to an appropriate amount of haptic guidance throughout training. Long term training experiments were conducted for an under-actuated target-hitting manual control task. The experimental results compare the efficacy of the novel progressive haptic guidance to two common fixed-gain haptic guidance schemes and virtual practice. The results from a month-long training experiment indicate that the proposed progressive shared control scheme reduces guidance interference as compared to fixed-gain guidance schemes thus increasing training efficacy.

Index Terms: H.1.2 [Model and Principles]: User/Machine Systems, – human factors— [H.5.2]: Information Interfaces and Presentation—User Interfaces - Haptic I/O, theory and methods, evaluation/methodology

1 INTRODUCTION

Whenever haptic guidance schemes are active, be they virtual fixtures, record and replay methods, or shared control paradigms, the participants experience task dynamics that are altered from the real task. Consequently, the assisted task serves as a secondary task to be learned. If the differences between primary and secondary tasks are large, severe interference can be experienced by the participants [5, 19, 20]. Most training schemes with haptic guidance utilize fixed-gain assistance that does not depend on the participant's performance. As a result, the participant might become dependent on the assistance in order to successfully complete the task and thus may not develop the required motor skills for the unassisted task.

The most common form of haptic guidance is achieved through the introduction of perceptual constraints on the workspace via so called *virtual fixtures* [18]. Virtual fixtures, generally employed as haptic guidance for performance enhancement in virtual environments [3, 8], are shown to be ineffective for training, since trainees tend to become dependent on the existence of the virtual fixtures to complete the task [16]. Dead zones are widely implemented to reduce the participant's dependence on the guidance. Analogous to using training wheels when learning to ride a bicycle, virtual fixtures with dead zones introduce forbidden regions to the task space, and the haptic guidance becomes available intermittently only to prevent large or unsafe errors [4]. Even though such methods provide a solution to the problem of the participants' dependence on the guidance, from the perspective of training, this kind of assistance provides nothing more than a safe medium for practice. Hence, the guidance provided by virtual fixtures with dead zones is not intended to assist the mechanism of learning because learning still takes place through *virtual practice*.

Another form of haptic guidance is motivated through teaching by demonstration. In these *record and play* strategies [6, 7, 9, 10, 21], the dynamics of an expert are recorded while performing the task, and are then played back to the novice to assist learning. The record and replay training scheme does not account for differences due to user-specific dynamics, and also prevents the novice from forming his or her own strategy for completing the task. Results from studies on record and replay effectiveness for motor skill training are highly inconclusive [6,7,9,10,21]. To increase effectiveness of record and play strategies, a progressive 4-step training scheme is proposed by Bayart *et al.* [1] that mimics the four stages of training for a bicycle riding task. These stages include demonstration, assistance from both trainer and training wheels, a "training wheels only" stage, and finally practice without training wheels. Bayart *et al.* tested this approach for a 3-D maze task and demonstrated the positive efficacy of such an approach.

Other performance-based progressive training schemes have been proposed as a way to gradually reduce the amount of guidance during training. Bell *et al.* proposed a performance-based progressive guidance scheme for self-learning of a computer-based radar-tracking simulation task, which showed significant beneficial effects [2]. A performance-based progressive robot-assisted therapy for stroke patients was proposed by Krebs *et al.* [11] in the field of rehabilitation, one of the major applications for haptic training. In Krebs' approach, the patients were provided with guidance during a reaching task by means of a virtual spring pulling them towards the target. The spring coefficient, hence the amount of guidance, was dependent on performance of the patients. Unfortunately, no conclusive training results are reported for this study. Similarly, in another robot-assisted rehabilitation study for gait training, human motor adaptation to dynamic environments was modeled as an error corrective learning process and the control gains of the guidance robot were adjusted at each trial based on the error [17]. The results from the interaction simulations of this study suggest that providing guidance only when needed is more effective than always assisting with a fixed amount.

O'Malley *et al.* [15] have proposed *shared control* as the most general active haptic guidance scheme for training. A shared controller dynamically intervenes, through an automatic feedback controller acting upon the system, to modify the (coupled) system dynamics during guidance. Li *et al.* showed that exposure to a *fixed-gain* error-reducing shared controller had a detrimental effect on participant performance of the target-hitting task at the conclusion of a month-long training protocol, regardless of the duration of exposure to the shared controller [13]. Li *et al.* hypothesized that the negative efficacy is mainly due to the gains of the controller being fixed thereby allowing participants to become dependent on the guidance during successful task completion.

In this paper, we demonstrate that a *progressive* shared control guidance scheme reduces the dependency of participants on the guidance by adjusting the control gains based on individual participant performance. In other words, a progressive shared control algorithm exposes participants to an appropriate amount of haptic

*e-mail: yvonneli@rice.edu
†e-mail: jhuegel@rice.edu
‡e-mail: vpatoglu@sabanciuniv.edu
§e-mail: omalleym@rice.edu.com

guidance based on their performance. Our results show that the progressive shared control protocol provides a significant increase in performance as compared to two other fixed-gain guidance protocols, namely shared control and virtual fixtures. However, a comparison between the progressive shared control scheme and a reference virtual practice (no guidance) scheme did not exhibit differences in performance.

This paper is organized as follows: Section 2 presents the methods used including the task, participants, performance measures, haptic guidance design, experiment design and procedure as well as data analysis. Section 3 presents the results, Section 4 discusses the findings and contributions, and Section 5 draws the conclusions of this experiment.

2 METHODS

A long-term human subject experiment was conducted to investigate the efficacy of the proposed performance-based progressive shared control guidance scheme to reduce interference. This training protocol aims to provide guidance at the beginning of each training session, then gradually adjust the amount of guidance based on the participant's performance, to ideally approach virtual practice by the end of training. The progressive shared control guidance is compared to two fixed-gain guidance schemes, namely shared control and virtual fixtures, as well as a virtual practice control group during training in a manual control task.

2.1 Task

The task for the training experiment is a target-hitting manual control task depicted in Fig. 1. Participants view the virtual double-mass spring system on a computer monitor and are asked to control the motion of mass m_1 via a two degree-of-freedom haptic device, a joystick. Through the two-mass system's dynamics, the participants are able to indirectly control mass m_2 to alternately hit a pair of fixed targets. 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 [16]. Figure 2 illustrates a participant sitting in front of the haptic interface system with the virtual environment displayed on the monitor. The virtual environment display includes a pair of targets and the double mass system. At any given time, one target is active, indicated by a changed color. The other is the inactive target. After m_2 contacts the active target, the target colors switch to indicate that the previous inactive target is now active. The targets are equidistant from the origin; therefore, the participants need to move the joystick, directly coupled to m_1, rhythmically, along the sloped path (referred to as the target axis), to cause m_2 to alternately hit the target pair. The task objective, as presented to each participant, is to hit as many targets as possible in each 20 second trial.

2.2 Participants

Thirty-two participants (8 females, 24 males, ages 18–33, 1 left-handed), primarily undergraduate students in engineering, participated in the experiment. The participants were given extra credit for an engineering class upon completion of the training. All participants signed consent forms approved by the IRB of Rice University to allow human performance data to be obtained and analyzed.

2.3 Performance Measures

Two performance measures were analyzed to assess participant performance for the target-hitting task, namely normalized hit count and average error. Normalized hit count is the total number of target hits within one trial normalized by the natural frequency of the corresponding dynamic system. Average error is the average of the instantaneous trajectory errors of mass m_2 as depicted schematically in Fig. 1. Together, these performance measures capture the features of the task, where normalized hit count gives an assessment of speed of execution, while average error monitors the ability of the participant to maintain a trajectory along the target axis. Average error is treated as a secondary performance metric, since participants are not specifically instructed to reduce the deviation of m_2 from the target axis. Nevertheless, these two measures provide a means for the proposed haptic guidance schemes to be objectively compared.

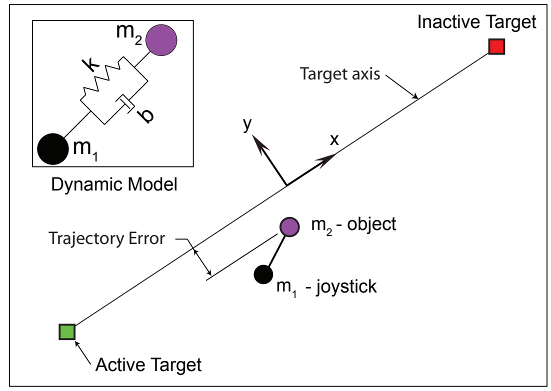

Figure 1: Target hitting task: the participant controls the location of m_1 (force feedback joystick) in order to cause m_2 (object) to hit the desired target. Inset shows virtual underactuated system. Trajectory error is defined as the deviation of m_2 (object) from the target axis.

Figure 2: A participant seated at the training station holding the force feedback joystick and interacting with the target-hitting task on the display.

2.4 Haptic Guidance

The four different haptic guidance schemes presented and subsequently compared in this experiment include progressive (condition P) and fixed gain shared control (condition S), as well as virtual fixtures (condition V) and virtual practice akin to no assistance (condition N). In the virtual practice (N) interaction mode, participants felt the forces generated solely due to the internal dynamics of the 2 DOF system. In contrast, for the virtual fixtures (condition V) and shared control cases (conditions P and S), participants felt the forces due to both the internal dynamics of the system and the guidance forces intended to assist but that may actually interfere in task completion. In the virtual fixtures (condition V) guidance mode, virtual walls were used to encourage users in a passive manner to move mass m_1 along the target axis, thereby causing m_2 to settle along the same path. The virtual walls generated forces proportional to the deviation and velocity of mass m_1 normal to the x-axis (see Fig. 1). The force for one virtual wall aligned with the x-axis was calculated according to

$$F_{py} = k_{wall}(y_1 - y_{wall}) + b_{wall}\dot{y} \tag{1}$$

These forces were subsequently displayed to the participant via the 2-DOF haptic joystick. The virtual wall parameters were chosen as $k_{wall} = 22.8N/m$ and $b_{wall} = 0.57\ Ns/m$. In the error reduction implementation of shared control (conditions P and S), the dynamics of the (state dependent) shared controller are designed such that the coupled (closed loop) dynamics of the system are simpler to manipulate than the system dynamics without the controller in place. Hence, by simplifying the task dynamics through coupling, the shared controller helps the participant to achieve better task performance. For the target-hitting task used in this paper, the forces to be displayed due to the error reducing shared controller, F_{sx} and F_{sy} can be expressed as

$$F_{sx} = 0 \tag{2}$$

$$F_{sy} = m_1\ddot{y}_1 - m_2\left[(K_v + 2\lambda)\dot{y}_2 + (K_p + \lambda^2)y_2\right] \tag{3}$$

Fixed control parameters used to implement the shared control guidance (condition S) were selected as $\lambda = 1rad/s$, $K_p = 70N/m$, and $K_v = 1Ns/m$. The detailed implementation of error reducing shared controller is described in the authors' previous work [13].

A performance-based progressive algorithm was employed to determine the gain (K_p) of the controller during guidance sub-sessions for progressive shared control (condition P). The input performance measurement for the algorithm was the normalized hit count since it is the primary goal of the task. The progressive shared control gain update law was controlled by a rolling average of three consecutive trials. Once the average of the current trial and two previous trials (average 2) is larger than the average of previous three trials (average 1), the control gain decreases. On the other hand, if average 2 is smaller than average 1 for three consecutive trials, control gain increases. Furthermore, if the absolute value of average 1 is above a certain threshold, 30 normalized hit counts, the control gain decreases. The control gain was adjusted based on the ratio of the difference between average 1 and average 2 over average 1. This update law, similar to the one-up three-down scheme described by Levitt [12], aims to decrease the haptic guidance, thereby decreasing the dependence on the guidance while the participant's performance still increases. In this way the progressive shared control scheme approaches virtual practice toward the end of training. Figure 3 shows the decaying progressive shared control gain of a typical participant. The control gain of the progressive shared controller, K_p, starts from the same value as the fixed-gain shared controller, and, depending on the performance of the participant, adjusts after each guidance trial throughout the experiment.

2.5 Experiment Design

The experiment was composed of 11 sessions, including an evaluation session, nine training sessions, and a retention session. Each training session contained three sub-sessions: pre-guidance baseline, guidance, and post-guidance baseline. Each sub-session consisted of 14 trials, with each trial lasting 20 seconds. Details of the experiment design are schematically represented in Fig. 4. The control group, also referred to as the virtual practice (N) group, received no haptic guidance during guidance sub-sessions of the experiment. However, this group did receive haptic and visual feedback of the task and environment as did all groups in all trials. The virtual fixtures (V) group received virtual fixture guidance during the guidance sub-sessions. The shared control (S) group was provided with guidance via the fixed-gain error reducing shared con-

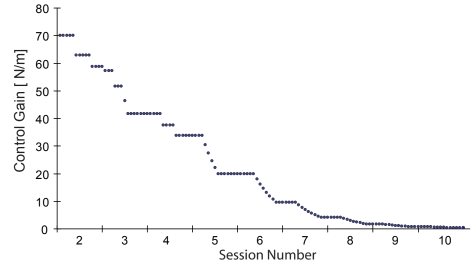

Figure 3: Progressive shared control gain (K_p) of a participant during guidance sub-sessions #2-10 illustrates the typical decaying trend. Decreasing gain is indicative of improving performance in terms of normalized hit count.

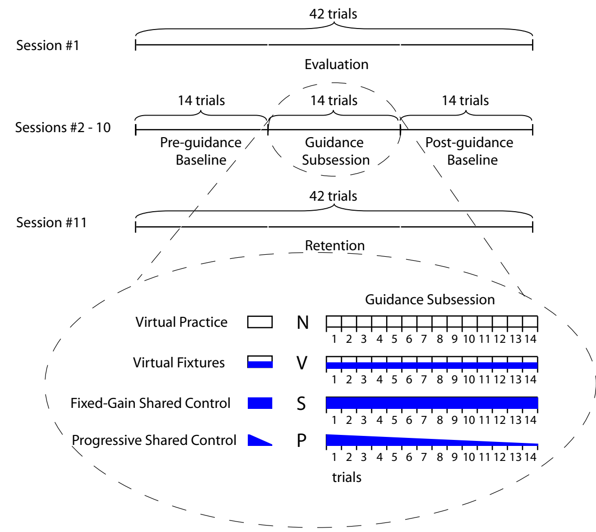

Figure 4: Schematic representation of the experiment design. The experiment consists of one evaluation, nine training and one retention session. Each training session contains three sub-sessions: pre-guidance baseline, guidance, and post-guidance baseline. During each guidance sub-session the virtual practice (N) group receives no haptic guidance while the virtual fixtures (V) group feels virtual fixture guidance for all 14 trials of the guidance sub-session. The shared control (S) group is provided with shared control guidance throughout all 14 trials of the guidance sub-session while the progressive shared control (P) group perceives different amounts of shared control throughout the guidance sub-session based on their individual performance

troller implementation. Finally, the performance-based progressive shared control (P) group received shared control guidance based on the current performance of the trainee. Pre-guidance and post-guidance baseline sub-sessions without the influence of haptic guidance were long enough to analyze dynamic interference effects which is the aim of this study.

Three sets of parameter values for the under actuated system were utilized as sub-tasks to increase the difficulty of the target-hitting task. Table 1 lists the three selected sets of system parameters that govern the dynamic response of this system. These parameter sets were varied in a controlled manner during the experiment to increase the complexity of the task, yet still enable data analysis and comparisons between groups, participants, and experiment sessions. Within each 14 trial sub-session, five repetitions of parameter sets 1 and 2, and four repetitions of parameter set 3 were

Table 1: Parameters of the two-mass spring damper system presented to the participants in order to increase task difficulty.

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

presented. The order of presentation was controlled in such a way that the first three trials of every sub-session contained one presentation of each set of system parameters. Similarly, the last three trials of every subsession contained one presentation of each set of system parameters.

2.6 Procedure

The participants were assigned to one of four training protocols based on their initial performance of the target-hitting task. Before the experiment, each participant was given a maximum of five minutes to become familiar with the haptic joystick and the virtual task. In order to control for individual differences in performance, each participant was asked to perform the task in an evaluation session, administered without haptic guidance. The purpose of the evaluation session was to measure the initial task performance of each participant so that well-balanced group assignments could be made. After the evaluation session, each participant was scored based on the total number of target hits. Participants were ranked according to their normalized hit count score, and were divided into eight sets with respect to their ranking. Then, equal numbers of participants from each set were randomly assigned into one of the four schemes (progressive shared control (P), fixed-gain shared control (S), fixed-gain virtual fixture (V), and virtual practice (N)) such that the average scores for the four groups were roughly equivalent at the start of nine training sessions.

All groups completed one evaluation session, nine training sessions, and one retention session. The virtual practice (N) group served as the control set with no haptic guidance provided during the guidance sub-sessions of the protocol. In order to assess the guidance interference and improvement of all participants across the nine training sessions, baseline sub-sessions of 14 trials without guidance were administered before and after each guidance sub-session. One guidance sub-session and its corresponding pre- and post-guidance baseline sub-sessions took place in one 30 minute session. The nine training sessions were separated by two to three days, such that the participants completed all the sessions in no less than three but no more than four weeks. One month after the final training session, all participants were recalled to complete one retention session. In the evaluation and retention sessions, no haptic guidance was provided to any participants.

2.7 Data Analysis

Repeated measure ANOVAs were utilized to determine significance of results. The experiment consisted of two factors, namely guidance mode and session. The guidance mode was between-subjects, with levels progressive shared control (P), fixed-gain shared control (S), virtual fixtures (V), and virtual practice (N). The session factor was within-subjects, with levels of evaluation, training (9 in all), and retention, for a total of 11 levels. In order to further explore the influence of different training schemes, another statistical analysis method, difference of least square means, was used. This analysis method takes into account all different groups by using an adjusted mean for each group that isolates the effect of each individual group, then gives out specific comparisons between each two groups.

3 Results

A month-long training experiment was conducted to investigate the effect of the proposed progressive shared control haptic guidance scheme to reduce guidance interference during training. Four conditions were compared, namely progressive shared control (P), fixed-gain shared control (S), fixed-gain virtual fixtures (V), and virtual practice (N). Figure 5 shows the performance for all four groups in the post-guidance baseline sub-sessions. The performance of all groups improved significantly in terms of both performance measures (normalized hit count and average error) as expected and saturated near the end of training. A repeated measures ANOVA with between-subject factors (group as between-subject factor, session as within-subject factor) was carried out to determine significance of results for these four groups. The results revealed a significant main effect of group and session for the post-guidance sub-sessions in terms of both normalized hit count and average error. A summary of these ANOVA results is listed in Table 2.

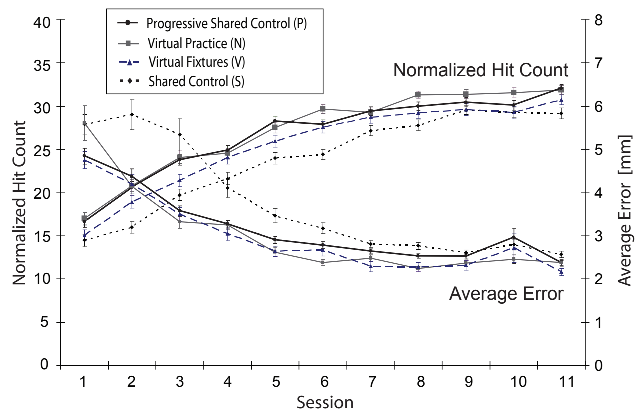

Figure 5: Post-guidance baseline normalized hit count and average error plots for different haptic guidance groups over eleven sessions of the training protocol.

Table 2: Summary of significance measured by ANOVA in terms of both normalized hit count and average error.

Measure	Effect	Post-Guidance
Normalized Hit Count	Group	$F(3,444) = 13.29,\ p < 0.0001^*$
	Session	$F(8,437) = 182.58,\ p < 0.0001^*$
	Interaction	$F(24,1268) = 2.28,\ p = 0.0004^*$
Average Error	Group	$F(3,444) = 19.67,\ p < 0.0001^*$
	Session	$F(8,437) = 67.16,\ p < 0.0001^*$
	Interaction	$F(24,1268) = 3.33,\ p < 0.0001^*$

A summary of all pertinent comparisons for the least square means analysis is listed in Table 3. The effect of group, as can be seen in Table 2, shows that the performance of the progressive shared control (P) group is significantly different from both the virtual fixtures (V) group and the fixed-gain shared control (S) group. Furthermore, Fig. 5 clearly shows that the progressive shared control (P) group outperforms both of the other forms of haptic guidance in terms of the normalized hit count which is objective measure of the task. The normalized total guidance force provided during the guidance sub-session is introduced as a way to investigate how much the participants depend on the existence of the haptic guidance. The normalized measure for each guidance sub-session is calculated as the percentage change of the total guidance force

Table 3: Summary of Significance Measured by Differences of Least Square Means in Terms of Normalized Hit Count and Average Error in the post-guidance sub-sessions.

Compare	Normalized Hit count	Average Error
P vs. S	$p < 0.001^*$	$p < 0.001^*$
P vs. V	$p = 0.003^*$	$p = 0.042^*$
P vs. N	$p > 0.05$	$p > 0.05$
S vs. V	$p < 0.001^*$	$p < 0.001^*$
S vs. N	$p < 0.001^*$	$p < 0.001^*$
V vs. N	$p < 0.001^*$	$p > 0.05$

provided within the session, compared to the total guidance force provided during the first session. This measure, as shown in Fig. 6, indicates the amount of guidance provided by each of the three haptic guidance schemes. Even though the reduced average errors for all groups (shown in Fig. 5) indicate that learning takes place for all training schemes, it is clear from Fig. 6 that only the assistance forces provided by the progressive shared control (P) group have a decreasing trend throughout training.

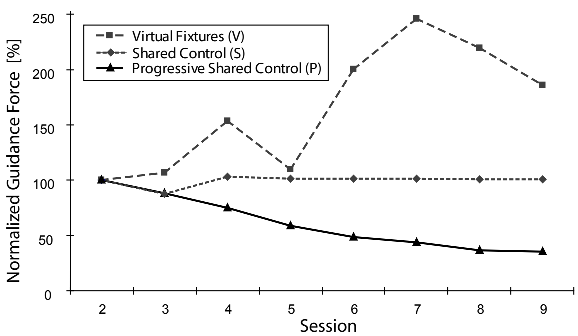

Figure 6: Total guidance force during the guidance sub-sessions for all three haptic guidance schemes demonstrates that the progressive shared control (P) group depended decreasingly on the guidance while the other two groups continued to depend on the guidance throughout training. Results are normalized based on the amount of guidance force incurred in the first session of training.

4 Discussion

A performance-based progressive shared control algorithm was implemented to reduce the effects of guidance interference observed with haptic guidance schemes for training. The guidance scheme focused on providing sufficient haptic guidance at the beginning of training, followed by decreasing the strength of guidance progressively based on the performance of the participant. This technique not only reduced the interference effects of haptic guidance, but also reduced the possibility of developing dependence on the haptic guidance.

The experimental results indicate that the progressive shared control scheme exhibits positive training efficacy compared to the two fixed-gain schemes as shown in Fig. 5. Since the fixed-gain error-reducing shared controller used in this study was implemented to reduce the deviation of m_2 from the target axis, the same amount of guidance force experienced throughout all training sessions would indicate that the deviation from target axis did not improve. On the other hand, the progressive shared control (P) group receives a decreasing total guidance force throughout the training sessions as shown in Fig. 6. Indeed, the portion of the total force displayed to participants of the progressive shared control group decreases across training, since the guidance force was designed to decrease as performance improved. In contrast, the fixed-gain groups shows steady, and even increasing, trends in the portion of total force that is attributed to the haptic guidance algorithm. For these groups, despite increasing performance (see Fig. 5), the participants still rely on the guidance force to assist in task completion throughout training. Thus we conclude that the effects of interference have been minimized by adjusting the amount of guidance throughout training.

Even though the progressive shared control (P) scheme performs significantly better than either fixed-gain haptic guidance scheme (shared control and virtual fixtures) it did not perform significantly better than the reference virtual practice (N) group in terms of normalized hit count. The first reason for this result is that our experiment design which uses equal-length sub-sessions is optimized to study effects of interference. In this design, subjects in the virtual practice (N) group are exposed to the primary task longer than the other groups who experience a secondary task which includes haptic guidance. Another possible reason may be due to the design of the shared controller. Decreasing interference is a necessary, but not sufficient, condition for haptic guidance schemes to achieve the training efficacy of virtual practice. Other factors (not considered in this study) may also play an important role in the efficiency of haptic guidance schemes. For example, according to Lintern et al., during training, a task should be simplified only if the important perceptual invariants of the task are preserved [14]. In the current implementation, the shared controller assists position control orthogonal to target axis, but neglects the temporal aspect of the control task. Time series plots for the same trial and session for three typical participants shown in Fig. 7 illustrate that exciting the system near its resonant frequency may be another critical component of the task. The participant with low performance (Fig. 7(a)) has only 6 target hits. The plot demonstrates that the participant is unable to stay on the x-axis and to excite the system rhythmically. The participant with high performance (Fig. 7(c)) has 34 target hits and demonstrates the ability to stay on the x-axis and to excite the system rhythmically. The performance of a nominal participant shown in Fig. 7(b) shows good trajectory error performance but is not exciting the system near the resonant frequency, resulting in only 17 hits. A redesign of the progressive shared controller to capture all essential aspects of the task may lead to significantly better training performance than virtual practice.

We have demonstrated that the detrimental effects of interference due to the virtual task dynamics being altered by haptic guidance schemes can be remedied by the implementation of the performance based progressive guidance approach. The results indicate that the participants of the progressive shared control group outperform participants who received either of two fixed gain haptic guidance schemes. Among the three haptic guidance groups, the progressive shared control (P) group is shown to be the only group that experiences progressively decreasing guidance force. All the haptic guidance groups shown learning trends across training; however, since the progressive shared control (P) group experience decreasing guidance forces as time progresses, these participants can develop the necessary skills for the task without developing dependency on the guidance. As a consequence, the performance of progressive shared control (P) group is significantly better than both (S and V) groups which experience fixed-gain guidance throughout training.

5 Conclusions

This paper tested and verified the hypothesis that a *progressive* shared control based training scheme, which adjusts its control gains based on participants' performance, can reduce the possible dynamic interference and dependence of participants on haptic guidance. Participants trained to complete a virtual target-hitting

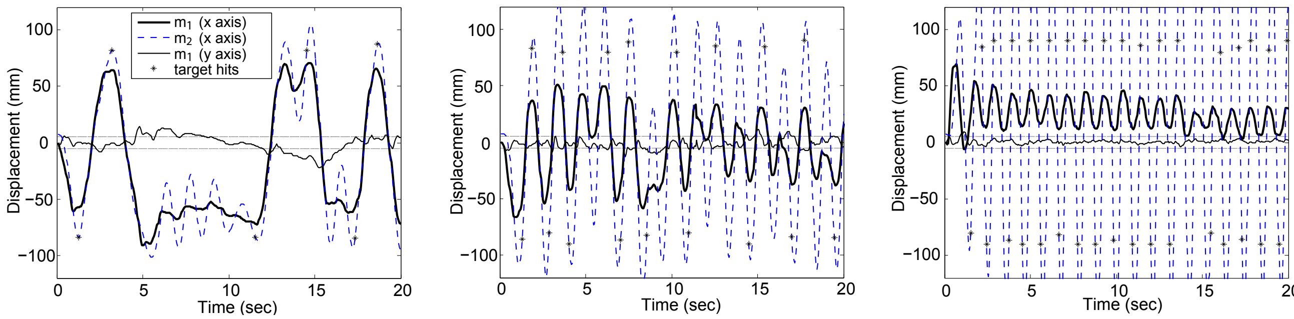

(a) low performance example (b) moderate performance example (c) high performance example

Figure 7: Displacement vs. time traces from trial 10 of session 4 for three typical participants. (a) shows the high average error and irregular input motion of a low performer. (b) shows the low average error but inconsistent input motion of an intermediate performer and (c) shows a high performer with very small average error and very consistent excitation. Results indicate that error-reduction alone may not be a sufficient strategy when trying to increase the number of target hits achieved

task with various types of haptic guidance and their performance was compared. The results indicate that the proposed progressive shared control scheme performs significantly better than fixed-gain haptic guidance schemes (fixed-gain shared control and fixed-gain virtual fixtures) yet not significantly better than virtual practice. The superior performance of the progressive shared control group over other haptic guidance groups is due to the performance-based progressive control algorithm that reduces the dynamic interference and dependence of participants on the haptic guidance and promotes acquisition of motor skills for the task. The approach is generalizable for haptic guidance based training of a range of complex tasks such as calligraphy, surgery and vehicle control.

ACKNOWLEDGMENTS

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

REFERENCES

[1] B. Bayart, A. Pocheville, and A. Kheddar. An adaptive haptic guidance software module for 1-touch: example through a handwriting teaching simulation and a 3d maze. In *IEEE International Workshop on Haptic Audio Visual Environments and their Applications*, pages 51–56, 2005.

[2] B. S. Bell and S. W. J. Kozlowski. Adaptive guidance: Enhancing self-regulation, knowledge, and performance in technology-based training. *Personnel Psychology*, 55:267–306, 2002.

[3] A. Bettini, P. Marayong, S. Lang, and A. M. Okamura. Vision assisted control for manipulation using virtual fixtures: Experiments at macro and micro scales. In *Proc. of the IEEE International Conference on Robotics and Automation*, pages 3354–3361, 2002.

[4] A. Bettini, P. Marayong, S. Lang, A. M. Okamura, and G. D. Hager. Vision-assisted control for manipulation using virtual fixtures. *IEEE Transactions on Robotics*, 20(6):953–966, 2004.

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

[6] D. Feygin, M. Keehner, and F. Tendick. Haptic guidance: Experimental evaluation of a haptic training method for perceptual motor skill. *IEEE International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 40–47, 2002.

[7] R. B. Gillespie, S. O'Modhrain, P. Tang, C. Pham, and D. Zaretsky. Virtual teacher. In *Proc. of the ASME Dynamic Systems and Control Division*, pages 171–178, 1998.

[8] D. P. Haanpaa and G. P. Roston. An advanced haptic system for improving man-machine interfaces. *Computers & Graphics*, 21(4):443–449, 1997.

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

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

[11] H. Krebs, J. Palazzolo1, L. Dipietro, M. Ferraro, J. Krol, K. Rannekleiv, B. Volpe, and N. Hogan. Rehabilitation robotics: Performance-based progressive robot-assisted therapy. *Autonomous Robots*, 15(1):7–20, 2003.

[12] H. Levitt. Transformed up-down methods in psycho-acoustics. *Journal of the Acoustical Society of America*, 49(2):467–477, 1970.

[13] Y. Li, V. Patoglu, and M. O'Malley. Negative efficacy of fixed gain error reducing shared control for training in virtual environments. *ACM Transaction on Applied Perception*, 6(1), 2009.

[14] G. Lintern. An informational perspective on skill transfer in human-machine systems. *Human Factors*, 33(3):251–266, 1991.

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

[16] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li. Shared control in haptic systems for performance enhancement and training. *ASME Journal of Dynamic Systems, Measurement and Control (JDSMC)*, 128(1):75–85, 2006.

[17] D. Reinkensmeyer, D. Aoyagi, J. Emken, J. Galvez, W. Ichinose, G. Kerdanyan, J. Nessler, S. Maneekobkunwong, B. Timoszyk, K. Vallance, R. Weber, J. Wynne, R. de Leon, J. Bobrow, S. Harkema, and V. Edgerton. Robotic gait training: Toward more natural movements and optimal training algorithms. In *Proc. of the 26th Annual International Conference of the IEEE EMBS*, pages 4819–4821, 2004.

[18] L. B. Rosenberg. Virtual fixtures: perceptual tools for telerobotic manipulation. In *Proc. of the IEEE Annual Virtual Reality International Symposium*, pages 76–82, 1993.

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

[20] E. M. Siipola and H. Israel. Habit interferences as dependent upon stage of training. *American Journal of Psychology*, 45:205–227, 1933.

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