Co-presentation of Force Cues for Skill Transfer via Shared-control Systems

Dane Powell* Marcia K. O'Malley†

Rice University

ABSTRACT

During training and rehabilitation with haptic devices, it is often necessary to simultaneously present force cues arising from different haptic models (such as guidance cues and environmental forces). Multiple force cues are typically summed to produce a single output force, which conveys only relative information about the original force cues and may not be useful to trainees. Two force co-presentation paradigms are proposed as potential solutions to this problem: temporal separation of force cues, where one type of force is overlaid with the other in staggered pulses, and spatial separation, where the forces are presented via multiple haptic devices. A generalized model for separating task and guidance forces in a virtual environment is also proposed. In a pilot study where sixteen participants were trained in a dynamic target-hitting task using these co-presentation paradigms, simple summation was in fact most effective at eliciting skill transfer in most respects. Spatial separation imposed the lowest overall workload on participants, however, and might thus be more appropriate than summation in tasks with other significant physical or mental demands. Temporal separation was relatively inferior at eliciting skill transfer, but it is hypothesized that this paradigm would have performed considerably better in a non-rhythmic task, and the need for further research is indicated.

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces—Haptic I/O

1 INTRODUCTION

One of the many benefits of training in haptic-enabled virtual environments is that movements can be monitored and shaped in real time via the use of perceptual overlays. For instance, Rosenberg proposed the use of virtual fixtures that passively prevent participants from entering forbidden regions of a work environment and could be used to constrain a novice learner's motions to an optimal trajectory [1]. Gillespie et al. proposed the use of a virtual teacher, a more active form of guidance that instructs novices to perform dynamic tasks by giving them shared control of a task with a virtual expert [2]. O'Malley et al. showed that such shared-control systems were as effective as virtual fixtures at facilitating skill transfer [3]. Any time that such perceptual overlays are used, there are two principal types of forces that a novice will experience. "Task" forces arise from interactions with the virtual environment, and might consist of reaction forces generated by collisions with virtual objects or dynamic forces arising from the manipulation of massive objects. "Guidance" forces are generated by the perceptual overlay, and might be used to constrain a novice to some optimal trajectory or prevent him or her from entering dangerous or forbidden regions of the environment. This presents a problem: how is it possible to display both of these types of forces to a user simultaneously?

*danep@rice.edu

†omalleym@rice.edu

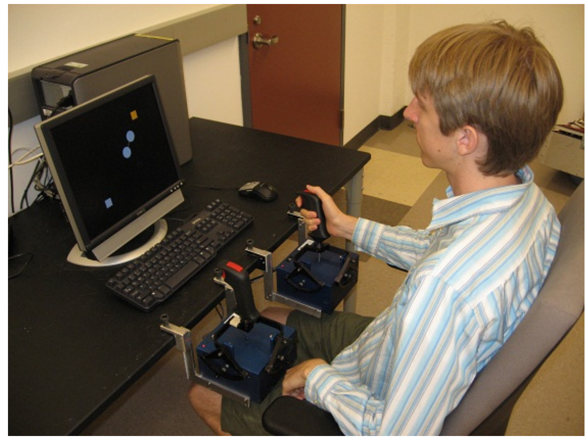

Figure 1: A participant performing an evaluation trial in a pilot study. Note that screen objects have been enlarged 10x for illustrative purposes.

The model of a virtual teacher proposed by Gillespie et al. replicates real-world teaching methods in order to facilitate skill transfer and reconcile this problem. He presents the example of a tennis expert teaching a novice how to swing a racket using hands-on demonstration. There are three ways that this demonstration could occur. In an "indirect contact" paradigm, the expert and the novice grasp the racket in different locations and perform the swing together. In a "double contact" paradigm, the novice grasps the racket while the expert grasps the novice's hand and guides the novice through the swing. In a "single contact" paradigm, the expert grasps the racket and the novice grasps the expert's hand. In the indirect and single contact paradigms, the task forces (those generated by the dynamics of the tennis racket) are simply summed with the guidance forces (those generated by the expert exerting control over the racket). In the double contact paradigm, the forces are separated spatially, with task forces being applied to the novice's palm and guidance forces to the back of his or her hand. Gillespie et al. hypothesized that this double contact paradigm would be the most effective at eliciting skill transfer, because it passes the greatest amount of haptic information to the novice and allows the novice to easily discriminate between guidance and task forces [2]. However, they were not able to conclusively determine whether this paradigm was indeed better than the others.

While replicating a real-world teacher is an elegant and intuitive approach to creating a virtual teacher, the utility of the double contact paradigm is limited to cases where the physical constraints of the task being taught allow for this specific type of spatial separation of forces. Presenting forces in this manner effectively requires haptic devices with up to twice as many degrees of actuation and significantly higher complexity. In some cases, presenting forces in this manner may simply not be possible given the physical constraints of the task. Providing guidance and task feedback via separate but *identical* haptic devices might be a more feasible solution.

This method of spatial separation is tested in this study.

The indirect and single contact paradigms are much simpler and less costly to implement from a physical standpoint, but lack the separation of forces provided by the double contact paradigm. The logical solution to this problem is try to separate the force cues temporally rather than spatially. Endo et al. proposed a method of temporal separation where a task force is periodically supplemented by a guidance force, and used this method to train participants to grip a virtual object using proper grasping forces and fingertip placements [4].

In a pilot study, these three co-presentation paradigms (summation, temporal separation, and spatial separation of forces) were implemented in a training environment and tested against each other as well as against a control condition as described in the following sections.

2 Methods

2.1 Task description

The target-hitting task used in this experiment was based largely on a task originally used by O'Malley and Gupta [3, 5]. Participants controlled the position of an on-screen pointer using a 2-DOF haptic joystick (Immersion, Inc.'s IE2000), as shown in Figure 1. This was connected to a 5 kg mass by a spring with 100 N/m stiffness, 3 Ns/m damping, and the equation of motion $F = m\ddot{x} + b\dot{x} + cx$, as shown in Figure 2. Thus, participants could control the position of the mass only indirectly. Two targets were positioned equidistant from the center of the screen and at a 45° angle to the horizontal. At any given time, one target was inactive (blue) and the other active (orange). The active target could only be "hit" by the swinging mass, at which point the opposite target would become active. Each task trial was 20 seconds long, and the general goal was to hit as many targets as possible in this time-frame. Thus, by moving the pointer at the resonant frequency of the system (0.71 Hz) along a straight line connecting the targets participants could achieve the highest hit-count possible (approximately 28 hits). The physics and haptics were rendered in C++ and updated at the servo rate of 1000 Hz, the visual display was rendered by OpenGL at a rate of 60 Hz, and experimental data was recorded at 100 Hz.

During training, participants shared control of this system with a virtual expert, represented on-screen by an orange pointer that tracked the optimal trajectory (a straight-line path between the targets at a frequency of 0.71 Hz). The participant and expert shared control of the system via a massless proxy. This proxy was positioned exactly half-way between the novice and expert at all times, and was in turn connected to the mass. Force information was transmitted unilaterally from the expert to the proxy - thus, novices' movements had no effect on the expert's trajectory. This setup allowed for the discrimination of guidance forces, which arose from the link between the participant and the proxy, and task forces, which arose from the link between the proxy and the mass.

2.2 Guidance methods

During training, each group of participants received haptic feedback in one of the following manners:

1. **No guidance** - Only task forces were displayed. Thus, participants could track the expert visually on-screen but received no haptic indication of his position. This served as the control condition.

2. **Summation** - Task forces and guidance forces were combined by simple summation and presented via a single joystick. The two types of forces were scaled so as to each have a peak magnitude of about half of the maximum force output level of the joystick.

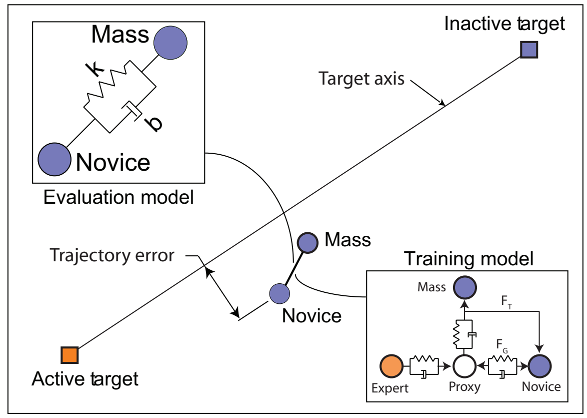

Figure 2: Diagram illustrating the task layout and dynamic models for evaluation and training trials. Participants control a pointer that is linked to a mass-spring-damper system and try to guide the mass to alternating targets. In evaluation trials, a simple spring and damper link the mass and the novice. In training trials, a series of "directional" spring and damper systems link the novice, expert, proxy, and mass, where arrows indicate the directions of force transfer. The expert follows the optimal trajectory irrespective of what is happening elsewhere in the system. The proxy's position is influenced equally by the expert and the novice. The novice receives guidance force F_G proportional to displacement from the proxy and task forces F_T proportional to the separation between the proxy and mass.

3. **Temporal separation** - Task forces were displayed at all times, and guidance forces were overlaid in 100 ms sinusoidal pulses at a frequency of 2 Hz (the optimal frequency and ratio as experimentally derived by Endo et al.). Participants described these guidance forces as "pulsating" and interpreted them as nudges or resistance that indicated the direction that they should be moving. The pulses were not frequent enough or large enough in magnitude to exert significant control over the task; thus, this mode prevented participants from becoming reliant on guidance forces, a problem described by Li et al. [6].

4. **Spatial separation** - Participants in this group used two joysticks during the experiment. Participants controlled the system using the primary joystick, onto which only task forces were displayed. Guidance forces were displayed on the secondary joystick so that its trajectory matched that of the expert's, also visible on-screen. Participants were instructed to lightly grasp this joystick with their non-dominant hand and to replicate the movements displayed there on the primary joystick. This allowed participants to intuitively mimic the expert's trajectory while still experiencing undistorted task dynamics. This paradigm also shares with temporal separation the advantage of forcing the participant to take control of the task and not rely on the expert to do any "heavy lifting".

2.3 Procedure

Participants performed the task over the course of 10 sessions on consecutive days, with each session consisting of 30 trials and thus taking 10-15 minutes. The experiment was broken up in this manner in order to combat fatigue and accelerate learning. This configuration of sessions was determined to elicit the fastest training in a pre-pilot study as compared to fewer sessions of greater length or more sessions of shorter length. Each session consisted of five evaluation trials ("pre-evaluation" trials), then twenty training trials, and finally five more evaluation trials ("post-evaluation" trials).

In evaluation trials, participants had sole control over the system via a single joystick and were instructed to hit as many targets as possible in the limited time-frame. The expert was not present or visible in any way, and thus the only haptic feedback was from the task dynamics. During training trials, the virtual expert and each participant shared control of the system under one of the experimental conditions, and participants were instructed to track the expert as closely as possible. In order to encourage participants to perform to the best of their ability and follow these instructions, gift cards were awarded to the participants that best achieved each of these goals. Participants were allowed a one-minute familiarization trial before the first session. During this trial, targets were an order of magnitude larger than in normal trials and were randomly located on the screen. This allowed participants to become familiar with the task without developing any significant task-specific skills.

At the end of each session participants also reported their perceived workload during the task by completing the NASA TLX questionnaire [7]. This questionnaire allows participants to rate their perceived workload on six different sub-scales: mental demand, physical demand, temporal demand, performance, effort, and frustration. It then lets them weight the contributions of each type of workload to the overall workload, and uses this information to compute a weighted average of the overall workload.

2.4 Data collection and analysis

For each trial, participants' performances were defined as their achieved hit count and average trajectory error. Trajectory error was defined as the mean absolute deviation from the straight-line path connecting each target. Thus, lower values of trajectory error indicate closer adherence to the optimal trajectory. Other performance metrics such as excitation frequency and work performed were considered but did not produce reliable results.

In order to analyze the data and compare performance between groups, each group's performance during evaluation trials was fit to the exponential learning model described below and given by Equation 1. Using this model rather than comparing group performance averages directly was necessary to overcome the inter-session variability of the data. This model also allowed for the prediction of future performance beyond participants' final sessions. This prediction capability was necessary because many participants did not reach their maximum performance potential, as evidenced by the positive slope in many of their learning curves at session ten and the fact that no participants approached the maximum theoretical hit count. Curves were fit both to post-evaluation trials in each session and to all evaluation trials in each session. Curves fit to pre-evaluation trials in each session did not produce reliable results. This is likely due to participants using these trials to reacquaint themselves with the task each day.

The coefficients of this exponential model each have a particular significance. The net amount of skill transfer is given by a. This can be thought of as the difference between a group's final potential performance level and their performance level before starting training. Thus, higher values of a indicate more skill transfer. The rate of learning is given by the time constant b. Lower values of b indicate that a group reached their maximum performance potential in a shorter period of time. This hypothetical maximum performance potential, which is the performance level that a participant would have reached given a large enough number of trials, is given by c.

$$y(t) = -ae^{-\frac{t}{b}} + c \quad (1)$$

2.5 Participants

A total of 17 participants enrolled in this study. One participant dropped out after his first session due to time constraints, and thus his data is not included. This left 16 participants evenly divided between the four experimental groups. All participants were males between the ages of 18 and 39 with no significant visual or motor impairments and no or little prior experience with virtual dynamic target-hitting tasks. Two participants were left-handed, and all participants controlled the task with their dominant or preferred hand. All participants provided their informed consent as approved by the Rice University Institutional Review Board.

3 Results

Confidence intervals were constructed for each of these curve-fit parameters for each group, as shown in Figures 3 and 4. The difference in parameter values in each potential pair of groups was tested for statistical significance using two-tailed t-tests at $\alpha = .05$. Statistically significant differences between groups were found when considering hit count and trajectory error as performance measures, as well as when comparing perceived workload.

3.1 Hit count

Considering post-evaluation trials alone, the temporal separation group showed reliably less skill transfer in terms of a than either the summation or spatial separation groups ($p = .006$ and $p = .013$, respectively). Results suggest that the summation group had a higher performance potential c than the spatial separation group ($p = .065$).

Considering all evaluation trials, the temporal separation group showed reliably less skill transfer in terms of a than the summation group ($p = .049$). The summation group also had a reliably higher performance potential c than the control or temporal separation groups ($p = .032$ and $p = .045$). Results suggest that the summation group learned faster (had a lower b) than the spatial separation group ($p = .077$).

3.2 Trajectory error

In post-evaluation trials, the temporal separation group showed reliably less skill transfer in terms of a than either the control or spatial separation groups ($p = .001$ and $p < .001$), and results suggest that temporal separation also produced less skill transfer than summation ($p = .06$). The summation group learned the task reliably faster than the spatial separation group ($p = .024$), and results suggest the summation group also learned faster than the control group ($p = .08$).

Considering all evaluation trials, all groups reduced their trajectory error faster than the spatial separation group ($p = .014$, $p = .002$, and $p = .002$). The summation group also had lower final trajectory error than the temporal separation group ($p = .022$).

3.3 Perceived workload

Groups trained using spatial separation reported a reliably lower overall workload than all other groups ($p = .004$ against control, $p < .001$ against summation, and $p = .026$ against temporal separation). They also reported reliably lower physical workload ($p < .001$ in all cases!) The control group reported reliably lower temporal demand than all other groups ($p = .021$, $p = .023$, and $p = .041$). The summation group reported reliably less frustration than any other group ($p < .001$, $p = .003$, and $p = .047$).

4 Discussion

In terms of hit count, results indicate that temporal separation elicits relatively little skill transfer and summation elicits a relatively high performance potential. These results are internally consistent when considering post-evaluation trials and all evaluation trials together. In terms of trajectory error, temporal separation elicited relatively little skill transfer, and spatial separation elicited relatively slow skill transfer, especially compared to summation.

The workload measures also produced some very compelling results. Spatial separation reduced the overall and physical workload experienced by participants during the task. This indicates that

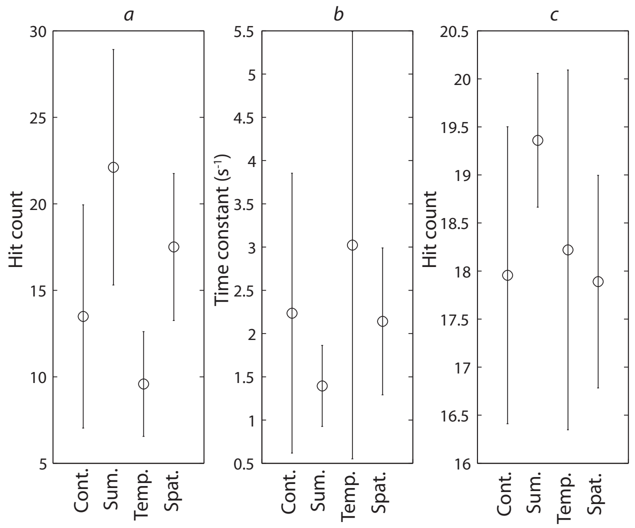

Figure 3: Hit count improvement amount, improvement rate, and final value by co-presentation paradigm in post-evaluation trials. Lines indicate 95% confidence intervals.

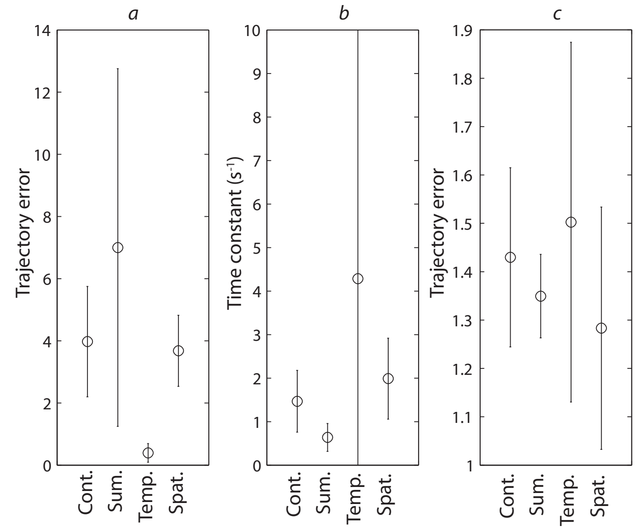

Figure 4: Trajectory error improvement amount, improvement rate, and final value by co-presentation paradigm in post-evaluation trials. Lines indicate 95% confidence intervals.

there might be a significant reason to prefer this paradigm in situations that demand physical and mental resources for other tasks, or in prolonged tasks that might lead to physical fatigue.

The fact that temporal separation seems inferior to other paradigms in almost all respects is most likely due to the rhythmic nature of the task. While there is an optimal excitation frequency and a clearly defined optimal path that minimizes trajectory error, the initial conditions of the task will produce optimal trajectories that are out of phase with each other in time. In other words, while it is true that following the expert precisely would elicit the highest hit count in the task, following the expert is not a necessary condition for achieving the highest hit count. It is entirely possible to follow the expert at a phase lag and still achieve the maximum hit count. In fact, in the summation condition, guidance forces and task forces are actually equal and opposite when the novice is out of phase with the expert by a certain amount. This might explain why the summation group reported the least amount of frustration during the task - they could actually complete training by being considerably more passive than other groups. By contrast, the temporal separation group reported that they found the assistance from the virtual expert to be pervasive and annoying even during later sessions, confirming that participants were likely performing the task out of phase with the expert. It follows that results might be very different in a non-rhythmic task. This indicates a significant need for further investigation into the relative merits of each of these co-presentation paradigms.

5 Conclusions

We have shown that for a rhythmic task, a simple summation paradigm is generally superior at co-presenting guidance and task forces, while temporal separation is generally inferior. However, we have also shown that the spatial separation paradigm might be preferred for tasks with high mental or physical demands. These are significant findings, as there are currently no published results directly comparing all three co-presentation paradigms for task training in a virtual environment. A generalized method of separating task forces from guidance forces in a shared-control system using a massless proxy was also proposed. Finally, we hypothesized that temporal separation would fare much better in a non-rhythmic task and established a need for further investigation.

Acknowledgements

This work was supported by NSF grants IIS-0448341 and IIS-0812569.

References

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

[2] R. B. Gillespie, M. O'Modhrain, P. Tang, D. Zaretzky, and C. Pham, "The virtual teacher," in *Proceedings of the ASME Dynamic Systems and Control Division*, vol. 64, 1998, pp. 171–178.

[3] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared control in haptic systems for performance enhancement and training," *Journal of Dynamic Systems, Measurement, and Control*, vol. 128, no. 1, pp. 75–85, Mar. 2006. [Online]. Available: http://link.aip.org/link/?JDS/128/75/1

[4] T. Endo, H. Kawasaki, K. Kigaku, and T. Mouri, "Transfer method of force information using Five-Fingered haptic interface robot," in *Proceedings of the Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE Computer Society, 2007, pp. 599–600. [Online]. Available: http://portal.acm.org/citation.cfm?id=1264885

[5] M. O'Malley and A. Gupta, "Skill transfer in a simulated underactuated dynamic task," in *Proceedings of the 2003 IEEE International Workshop on Robot and Human Interactive Communication*, 2003, p. 315–320.

[6] Y. Li, V. Patoglu, and M. K. O'Malley, "Negative efficacy of fixed gain error reducing shared control for training in virtual environments," *ACM Trans. Appl. Percept.*, vol. 6, no. 1, pp. 1–21, 2009. [Online]. Available: http://portal.acm.org/citation.cfm?id=1462058

[7] S. G. Hart, "Development of NASA-TLX (Task load index): Results of empirical and theoretical research," NASA Ames Research Center, Tech. Rep., 1988.