Efficacy of Shared-Control Guidance Paradigms for Robot-Mediated Training

Dane Powell[*] Marcia K. O'Malley[†]

Rice University

Abstract

Robot-mediated training has seen a plethora of implementations of shared-control haptic guidance, intended to teach novices to perform dynamic tasks by providing them with real-time visual and haptic feedback from real or virtual experts. The efficacies of these paradigms are difficult to quantify and compare, as the paradigms have typically been developed in an ad-hoc manner to suit specific devices and tasks. This work proposes a novel guidance paradigm taxonomy intended to help classify and compare the multitude of implementations in the literature, as well as a revised proxy rendering model to allow for the implementation of two relatively novel guidance paradigms (in addition to existing paradigms). The efficacies of these two paradigms, plus two more paradigms representing the vast majority of implementations in the literature, are compared in a controlled study with 50 healthy subjects. The results show that none of these paradigms are superior to visual-only guidance (the control condition), and that the two newer paradigms are actually detrimental to training for a target-hitting task. These results contradict many intuitions about how haptic guidance should be implemented and raise doubts about the efficacies of the most commonly-implemented paradigms.

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

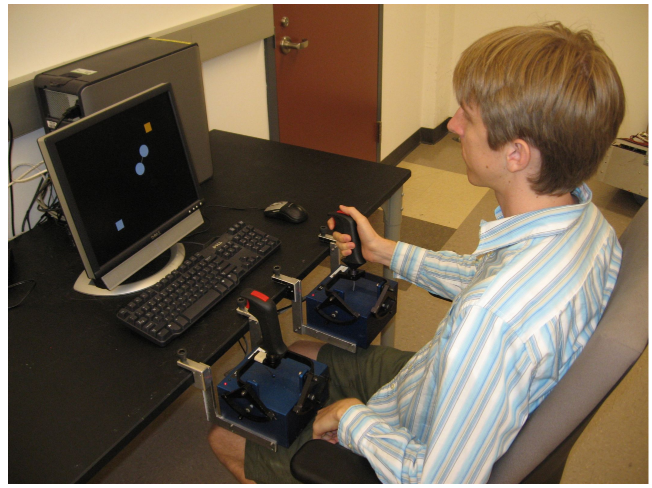

Figure 1: A subject performing a target-hitting dynamic task.

1 Introduction

Dynamic tasks are part of our everyday lives. Shooting a basketball, driving a car, or simply taking a sip of water are all characteristically dynamic tasks that require sensory feedback (especially haptic feedback), on-line movement planning, and adaptation to changing task conditions. Most importantly, these are all tasks that have one or more optimal solutions that either maximize a "positive" metric, such as likelihood of making a basket, or minimize a "negative" metric, such as the amount of effort required. These optimal solutions are learned through a combination of practice and training, either by direct intervention from a coach or through focused observation of other people performing the task. Similarly, there are many less common but more consequential dynamic tasks requiring extensive training, such as performing a laparoscopic surgery, flying an airplane, or teleoperating a remotely-operated vehicle.

Training for these tasks can be either human-mediated or robot-mediated. An expert surgeon gripping a novice's hand in order to physically help that novice complete a surgery would be an example of human-mediated training. Conversely, a novice training to complete the surgery in a virtual environment with the assistance of either a live or virtual expert surgeon would be an example of robot-mediated training.

While the question of how to apportion control of the system between expert and novice has been studied to some extent in the literature, the question of how to provide haptic feedback, especially to the novice, has been studied comparatively little. While haptic feedback can enhance a novice's sense of presence and cooperation [1, 2], its efficacy at improving training outcomes has not been thoroughly demonstrated. Such feedback, if coming directly from an expert (be they human or pre-programmed), is generally referred to as haptic "guidance," as it is generally used to guide a novice through the successful completion of a task.

Most guidance schemes used for robot-mediated training have been developed in an ad-hoc fashion to work with a specific device or task, making it difficult to compare the multitude of guidance schemes present in the literature. We propose that the many various existing guidance schemes can be distilled into a set of essential and representative characteristics, and that these characteristics can be used to develop a taxonomy for classifying guidance paradigms, as discussed in Section 3. Some of the more novel paradigms cannot be easily implemented using the traditional proxy rendering model; thus, an improved shared-control proxy model is proposed in Section 4. The efficacies of four guidance paradigms at training 50 healthy subjects to perform a target-hitting task are compared in a controlled study as described in Section 5. The results presented in Section 6 demonstrate that none of these haptic guidance paradigms are significantly better than visual guidance alone, and in fact some are significantly worse. The implications of these findings are discussed in Section 7.

2 Background

Guidance during robot-mediated training is usually provided via simple perceptual overlays such as virtual fixtures. Virtual fixtures, as proposed by Rosenberg [3], are simply perceptual overlays that passively prevent participants from entering forbidden regions of a work environment, and are most often used to constrain a novice's motions to an optimal trajectory.

[*]danep@rice.edu

[†]omalleym@rice.edu

Guidance might also take a more active form, such as the "record-and-replay" strategy used by Gillespie et al. [4] to train novices to balance a inverted pendulum. Such "assistive" methods are based on a number of intuitions about how people learn to perform visuo-motor tasks. Unfortunately, there is little evidence to back up some of these intuitions or to suggest how they can best be applied to enhance the efficacy of assistive strategies. A common assumption is that physically guiding a novice through the successful completion of a task will help the novice to somehow internalize and encode that pattern, and thus help the novice to repeat the pattern on his or her own in the future. While sounding plausible, this assumption is only weakly supported by the literature in the context of rehabilitation [5, 6], and has been refuted in many cases in the context of training healthy individuals [7, 8, 9]. Schmidt and Bjork [10] showed that guidance in many sorts of training (not just in visuo-motor tasks) can actually impair learning and retention, and proposed the "guidance hypothesis" to account for this discrepancy between the expected and actual results of guidance-based training.

The probable flaw in the assumption that assistive guidance improves training is that while the proprioceptive sensory pathways are active in the presence of guidance, the motor pathways are comparatively less active. Israel et al. [11] showed that when physically guided through a task, novices tend to become "passive participants" and exert less energy (reflecting less motor pathway activity) than when they perform the task on their own. Shadmehr and Mussa-Ivaldi [12] showed that the CNS relies on encoding and storing control loops between proprioceptive input and motor output in order to perform dynamic tasks, and thus if this control loop is weak or absent in the presence of guidance, the CNS will not be able to encode and retain it as it would during practice.

Another problem with assistive guidance is that because novices are passive and constrained to an optimal trajectory, they are going to make fewer errors than they would during practice. Thoroughman and Shadmehr [13] and others have shown that error drives the learning of dynamic tasks and building of internal models, and thus assistive guidance is likely to impair learning by preventing the commission of error.

Finally, a significant problem with assistive guidance is that it corrupts the inherent dynamics of a task as perceived by the novice. Most guidance methods are impedance-based, meaning that they apply a force in order to control the novice's position. Thus, a movement made during practice will result in force-feedback based on the inherent task dynamics, while an identical movement during training will result in force-feedback based on some combination of the task dynamics and guidance forces. If novices spend a bulk of their time in training, then in effect they will be learning the wrong task, as found by Crespo and Reinkensmeyer [14], who say that "subjects who trained with guidance reacted as if the assistance provided on assisted trials was a perturbation rather than following its example."

Gillespie et al. [4] proposed the use of a virtual teacher, a more active form of guidance than virtual fixtures that instructs novices to perform dynamic tasks by giving them shared control of a task with a virtual expert. O'Malley et al. [15] showed that such shared-control systems were as effective as virtual fixtures at facilitating skill transfer. The model of a virtual teacher proposed by Gillespie et al. replicates real-world teaching methods in order to facilitate skill transfer and reconcile the problem of guidance force corrupting task dynamics. 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, as shown in Figure 2. 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. [4] 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. However, they were not able to conclusively determine whether the double contact paradigm was better than the others.

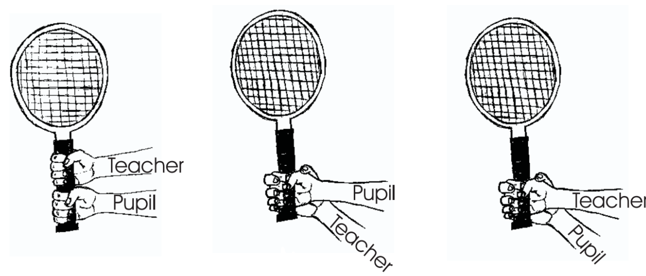

Figure 2: Gillespie et al. [4]'s Virtual Teacher paradigms. From left to right: indirect-contact, double-contact, and single-contact paradigms.

3 Guidance Paradigm Taxonomy

We propose that all guidance paradigms currently implemented in the literature in human-human, human-robot, and human-robot-human training architectures can be classified as one of the five types in this section based on two characterizing factors. By abstracting the principles of existing guidance paradigms from their specific implementations, we can develop a set of representative paradigms from the taxonomy and then compare the effectiveness of each of those paradigms while holding constant the specifics of the implementation (such as the choice of haptic device and dynamic task). The first factor that differentiates guidance paradigms is whether they assist or resist the novice in completing the task. The second factor is how paradigms reconcile the co-presentation of task and guidance forces, whether by separation or summation.

3.1 Gross Assistance

Classic virtual fixtures are the archetypal example of gross assistance (GA). By their nature, virtual fixtures have to be relatively stiff in order to keep novices from entering forbidden regions of the workspace, and thus guidance forces generated by collisions with virtual fixtures will dominate any extant task forces. Simple spring-damper couplings or attractor potential models used to "pull" novices towards a target are also typically implemented as GA, and can interfere with the perceived dynamics of tasks in a more subtle way than virtual fixtures. Shared-control guidance schemes such as the indirect-contact and single-contact virtual teacher paradigms also qualify as GA.

Gross assistance has been shown to be generally ineffective at improving training outcomes compared to practice without guidance. Reinkensmeyer [5] showed in simulation that "continual guidance" (GA) is "never beneficial compared to no assistance". Crespo and Reinkensmeyer [16] showed that "fixed guidance" (GA) produced only "slightly better immediate retention than did training without guidance," but did not show that this improvement was statistically significant.

3.2 Temporally Separated Assistance

The characterizing factor of temporally separated assistance (TSA) is that it separates guidance and task forces temporally, displaying each type alternately in quick succession via a single haptic device. Novices primarily experience unadulterated task forces, but training is frequently (on the order of 1 Hz) punctuated by brief periods of pure guidance, intended to "cue" novices as to the appropriate movements to make. In this way, the expert exerts "cognitive dominance" over the novice, while allowing the novice to retain "physical dominance"- in other words, allowing a novice to commit errors and actively generate movement plans in order to better learn the task dynamics.

Endo et al. [17] are the only group known to have proposed and tested a TSA paradigm. In a pilot study, they showed that TSA was effective at training subjects to grip a virtual object using proper grasping forces and fingertip placements. However, they did not study its effectiveness at training for dynamic tasks, and there are no other implementations of TSA in the literature.

3.3 Spatially Separated Assistance

Whereas TSA separates the presentation of task and guidance forces temporally in order to present them via a single haptic channel, spatially separated assistance (SSA) makes use of two haptic channels in order to present task and guidance forces simultaneously via the separate channels. The first and perhaps best example of SSA is the double-contact paradigm proposed by Gillespie et al. [4], which makes use of a specialized haptic device in order to present guidance from a virtual expert via one haptic channel (through the back of a novice's hand) and forces arising from the task dynamics via a second channel (through the novice's palm). Gillespie et al. [4] could not conclusively show that SSA was superior to practice.

There are no other known implementations of SSA in the literature, likely due to the relative complexity and propriety of the haptic devices necessary to implement e.g. the double-contact paradigm. While replicating a real-world teacher is an elegant and intuitive approach to implementing SSA, 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, and is tested in this study.

3.4 Gross Resistance

Gross resistance (GR) can take a number of different forms, but is generally characterized by increasing the difficulty of a task or resisting a novice's optimal completion of a task in some way. The theory behind GR is simply based on over-training: after training extensively in the presence of artificial resistance, novices will find it relatively easy to execute the same task in the absence of the resistance. There are three common implementations of GR: as a constant force-field or viscous force opposing movement, as a force that augments errors, or as forces producing random disturbances.

Lee and Choi [18] showed that training in the presence of random noise-based disturbance was superior to GA and practice at training healthy novices to perform a path-following task. Such noise-based GR has not been discussed elsewhere in the literature and is a prime candidate for further evaluation.

4 SHARED-CONTROL PROXY MODEL

Zilles and Salisbury [19] proposed a "constraint-based god-object" rendering algorithm (commonly referred to as a "proxy model") for calculating and displaying interactions between a haptic interface and a virtual environment. In this traditional proxy model, a massless "god-object", "avatar", or "proxy" represents the user in the virtual environment, and must obey all of the physical constraints of the virtual environment (i.e. walls, friction, etc...). The proxy is then connected to the haptic device by a virtual spring and damper coupling. This coupling allows the haptic device to penetrate virtual surfaces without necessarily leading to instability or requiring a specialized physical model.

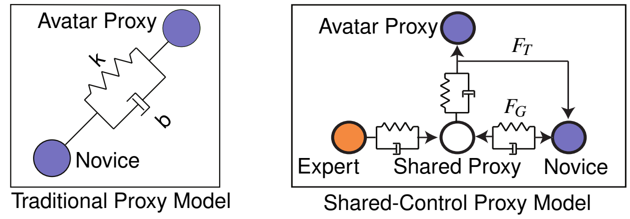

Figure 3: Traditional and Shared-Control Proxy Models. k is stiffness, b is damping, F_G is guidance force, and F_T is task force.

If a perceptual overlay or virtual expert is added to the environment, one can imagine that there are two qualitatively different types of forces in the system: "guidance" forces, which arise from interactions with the perceptual overlay or virtual expert, and "task" forces, which arise from interactions with the virtual environment. A distinction should be made between these types of feedback because they should contribute to a user's learning in fundamentally different ways: "guidance" forces should be used to shape the user's actions, whereas "task" forces should be incorporated into the user's internal model of the environment. The problem with the traditional proxy model is that it cannot discriminate between guidance and task forces in shared-control systems, and thus the forces are confounded when displayed to the user. This precludes the use of the more advanced separation guidance paradigms described in Section 3.

The proposed shared-control proxy model overcomes this deficiency by adding a second proxy and replacing the traditional spring-damper couplings with a series of "biased" spring and damper couplings. Whereas traditional couplings can only exert equal and opposite forces on attached nodes, biased couplings can exert opposite but arbitrarily scaled forces on each node and are only realizable in a virtual environment, as they essentially break Newton's Third Law. These couplings link the novice, expert, "shared proxy", and "avatar proxy" as illustrated in Figure 3, where arrows indicate the general directions of force transfer (in other words, the end of the coupling with a higher force gain).

The massless shared proxy's position is influenced equally by the expert and the novice, but is not influenced at all by the position of the avatar proxy, nor does it interact with the virtual environment. Thus, the shared proxy remains exactly between the novice and expert at all times, representing the averaged input of the novice and expert. Note that this average could be weighted in order to adjust the control authority α (as proposed by Nudehi et al. [20]) by simply changing the relative stiffnesses of the expert and novice couplings. The force generated by the coupling between the novice and shared proxy represents a pure guidance force F_G, since it is proportional to the deviation from the expert and unaffected by the virtual environment. Note that in this case, the expert will not receive any force feedback and thus will not be affected by the novice, which is the logical setup for tasks with a virtual expert. However, with a human expert present, force-feedback could be provided in a way similar to how the novice receives force feedback.

The avatar proxy must obey all of the constraints of the virtual environment and is coupled to the shared proxy, so that in free space

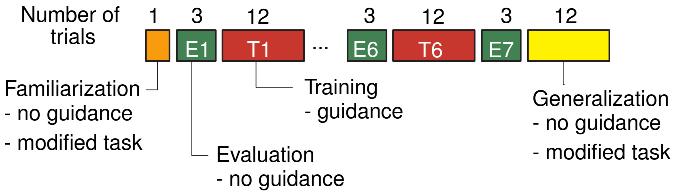

Figure 4: Illustration of session structure.

both proxies ideally share the same position. However, when the user comes into contact with a virtual surface, the invisible shared proxy will penetrate the surface to the same extent as the haptic device, while the avatar proxy will remain outside the surface. The force generated by the coupling between the two proxies then represents a pure task force F_T, since it is proportional to the deviation between the commanded position of the shared proxy and actual position of the avatar proxy.

5 Methods

5.1 Experimental Design

Four prototypical guidance schemes were implemented using Immersion IE2000 2-DOF haptic joysticks, and their effectiveness at training subjects to perform two dynamic tasks was evaluated in a 50-subject controlled study. Subjects trained with the assistance of a virtual expert using the shared-control proxy model described in Figure 3, which allowed for the discrimination of task and guidance forces. The physics and haptics were rendered in C++ and updated at the servo rate of 1000Hz, the visual display was rendered by OpenGL at 60Hz, and experimental data was recorded at 100Hz.

5.1.1 Evaluation and Training Trials

Subjects performed the tasks over a number of trials. Each trial was 20 seconds long and generally categorized as either an "evaluation" trial or a "training" trial. In evaluation trials, subjects had sole control over the system via a single joystick and were instructed to perform the task to the best of their ability. During training trials, a virtual expert was also present in the system. This expert followed a predefined optimal trajectory for each task, and shared control of the system with each subject under one of the experimental conditions. Subjects were instructed to track the expert as closely as possible during training; by exactly matching the expert, they could achieve the best score possible in each task.

5.1.2 Structure

Subjects performed the task over a single one-hour session consisting of 106 trials grouped into a number of different blocks, as shown in Figure 4. Evaluation blocks consisted of three evaluation trials, training blocks consisted of 12 training trials, and a final generalization block consisted of 12 generalization trials. Subjects were allowed a one-minute familiarization trial with an easier version of the task prior to starting the session, as well as a 5-minute break midway through the session in order to prevent fatigue.

5.2 Subjects

A total of 50 subjects enrolled in the primary study, and were divided evenly between 5 experimental groups: no guidance, GA, TSA, SSA, and GR. Five subjects were left-handed, 45 right-handed, 33 male, and 17 female. All subjects controlled the task with their dominant or preferred hand. All subjects provided their informed consent as approved by the Rice University Institutional Review Board, had no significant visual or motor impairments and no or little prior experience with virtual dynamic target-hitting tasks. Gift cards were awarded to the highest-performing subjects.

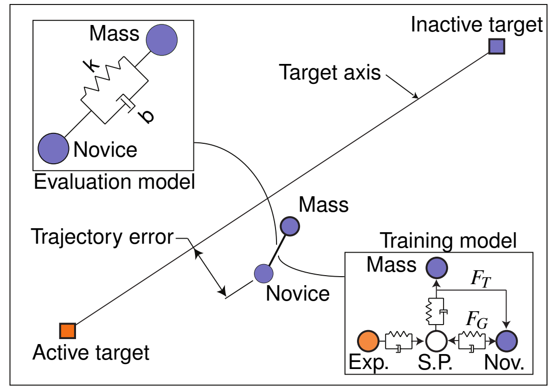

Figure 5: Target-hitting task and proxy models used for training and evaluation trials.

Table 1: Force outputs for guidance conditions. Both F_T and F_G are calculated according to the shared-control proxy model described in Section 4. The resistive force F_{PN} is calculated according to a Perlin noise function. For these experiments, $t_0 = 100\,\text{ms}$ and $t_1 = 500\,\text{ms}$.

Guidance	Force output (Joystick 1)
Control	$F_T(t)$
GA	$F_T(t) + F_G(t)$
TSA	$\begin{cases} F_T(t) + sin(\frac{t*\pi}{t_0})F_G(t) & \text{if } t \bmod t_1 \leq t_0; \\ F_T(t) & \text{if } t \bmod t_1 > t_0. \end{cases}$
SSA	$F_T(t)$
GR	$F_T(t) + F_{PN}(t)$

5.3 Task

The target-hitting task used in these experiments was based largely on a task originally used by O'Malley and Gupta [22] and O'Malley et al. [15]. Subjects controlled the position of an on-screen pointer using a 2-DOF haptic joystick (Immersion, Inc.'s IE2000), as shown in Figure 1. The joystick, displaying up to 9 N, was connected to a 5 kg virtual mass by a spring with stiffness $k = 100\,\text{N/m}$ and damping $b = 3\,\text{Ns/m}$, as shown in Figure 5. Thus, subjects 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 goal during evaluation trials 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, subjects could achieve the highest hit-count possible (approximately 28 hits). During training, subjects 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). During evaluation trials, subjects were instructed to hit as many targets as possible, while in training trials they were instructed to follow the expert as closely as possible.

5.4 Guidance Conditions

The mathematical representations of the guidance paradigms used during training trials are given in Table 1.

5.4.1 No Guidance (Control)

Only task forces were displayed as a control condition. Thus, subjects could track the expert visually on-screen but received no haptic indication of his position.

5.4.2 Gross Assistance (GA)

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.

5.4.3 Temporally Separated Assistance (TSA)

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. [17]). Subjects 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 subjects from becoming reliant on guidance forces, a problem described by Li et al. [21].

5.4.4 Spatially Separated Assistance (SSA)

Subjects used two joysticks during the experiment. Subjects 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. Subjects were instructed to lightly grasp this secondary joystick with their non-dominant hand and to replicate the movements displayed there on the primary joystick. This allowed subjects 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 subject to take control and do the "heavy lifting" in the task.

5.4.5 Gross Resistance (GR)

Task forces were combined with a randomly-generated disturbance force in the manner described by Lee and Choi [18]. A Perlin noise function with a nominal range of -1.2 N to 1.2 N was randomly generated for each joystick axis using the open-source *libnoise* library. At each timestep, the guidance force was generated from the values of these functions and summed with the task force to produce the net force displayed to the joystick.

6 Results

Outliers were defined for each cell (each unique combination of group and trial) as points further than 1.5 interquartile ranges from the cell mean and were replaced with their respective cell mean.

Mixed ANOVAs were performed on evaluation trial outcomes with hit counts as the dependent variable, guidance condition ("group") as a between-subjects factor, and trial number ("trial") as the within-subjects factor (repeated measure). Plots of the data for evaluation trials alone are shown at group-level in Figure 6. Outcomes for the omnibus ANOVA and for pairwise multiple comparisons, corrected using a Tukey-Kramer (TK) adjustment, are shown in Figure 7. A Ryan-Einot-Gabriel-Welch adjustment would preserve more power and be preferred, but was unavailable in the statistics software procedure used for analysis.

All groups exhibited a consisting learning trend. Multiple comparisons based on the mixed ANOVA showed that both the control and GR groups performed significantly better than the TSA and SSA groups, and that the GA group performed significantly better than the TSA group.

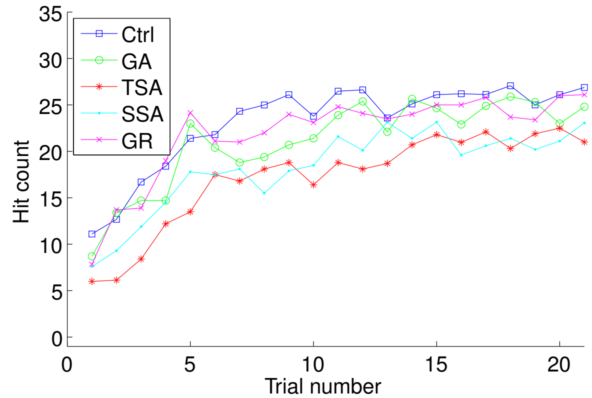

Figure 6: Mean group hit counts in evaluation trials, outliers replaced.

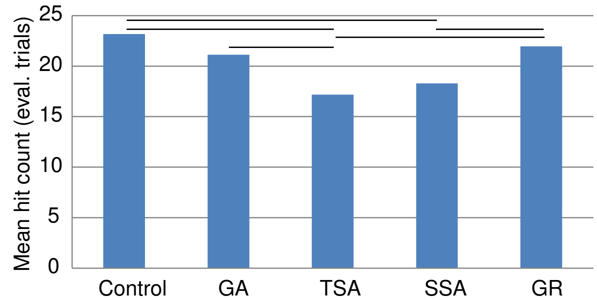

Figure 7: Mixed ANOVA results for evaluation trials. Fixed effect of group: $F(4,113) = 9.46$, $p < .001$. The interaction effect of group and trial was not significant. Lines indicate pairwise significance at $\alpha = .05$, family-wise error-corrected using a TK adjustment.

7 Discussion

Numerous previous studies have shown GA to be ineffective, as described in Section 3.1, so it is not surprising that it performed no better than the control group. Conversely, it is quite surprising that the separation paradigms (TSA, SSA) actually led to worse performance than either of the gross guidance paradigms (GA, GR), given that previous studies have shown TSA to be effective (as described in Section 3.2) and that SSA was specifically developed to overcome deficiencies identified with gross guidance. This unexpected result might be explained by several factors. Generally speaking, it is possible that subjects were simply unfamiliar with these novel forms of guidance, and did not fully understand how they were supposed to use the guidance. Anecdotally, subjects had a relatively easy time understanding the operating principles of GA and GR, while they had a comparatively harder time understanding how to use TSA and SSA. Thus, it is possible that more thorough training for how to use these somewhat complex paradigms would lead to improved results. Additionally, subjects were not informed of the theory behind the separation paradigms in order to obviate any placebo effects resulting from the power of suggestion or experimenter bias. However, it is possible that if subjects had understood *why* such complex guidance methodologies were being used, then they would be less frustrated with the guidance and better understand how to fully utilize it.

Overall, the results corroborate a strong form of the guidance hypothesis: namely, that any attempts at guidance (even resistive forms) can impair training as compared to practice. Previous studies have shown that subjects can become dependent on assistive guidance, and the guidance hypothesis theorizes that challenge is necessary to the learning process. The results of this study support that notion, but also suggest that even additional challenge can impair learning, and indeed that any interference with the task (through attempts at guidance) will impair training compared to straight practice. This is supported by the fact that both of the novel guidance separation paradigms (TSA, SSA), which were designed specifically to discourage dependency, led to significantly worse performance than the control group. Additionally, the results suggest that the gross guidance paradigms (GA, GR) led to worse performance than the control group, though this effect was not statistically significant.

These results indicate that perhaps a new approach to guidance is needed. For instance, instead of taking an "objective-oriented" approach and teaching subjects to simply follow an expert in order to complete task objectives, it might be more beneficial to take a "skill-oriented" approach to guidance by identifying and teaching the specific component skills necessary to complete a task. It is also possible that the best way to enhance training is to increase the difficulty of a task without altering the inherent task dynamics or interfering with task execution through explicit guidance. For instance, decreasing the target size might both be effective ways of enhancing training. Finally, this study compared combined visual and haptic guidance to visual-only guidance, and so the results might simply indicate that visual guidance is sufficient for complete training. This might be explained by the same reasoning behind the development of the separation guidance paradigms: by exerting cognitive dominance without any physical interference whatsoever, visual-only guidance is the ultimate separation guidance paradigm.

8 Conclusions

The results of this study show that many of our instincts about haptic guidance are wrong: conventional approaches to haptic robot-mediated training are not significantly better than practice, and more complex guidance paradigms can in fact be detrimental to the learning process. To facilitate continued research, this work has made a number of additional contributions. A guidance paradigm taxonomy has been proposed that will allow for easier discussion, classification, and comparison of haptic guidance paradigms. The traditional shared-control proxy model has also been improved in order to accommodate a number of more complex guidance paradigms, and a novel paradigm (SSA) has been developed.

References

[1] C. Ho, C. Basdogan, M. Slater, N. Durlach, and M. A. Srinivasan, "An experiment on the influence of haptic communication on the sense of being together," in *BT Presence Workshop*, 1998.

[2] I. Oakley, S. Brewster, and P. Gray, "Can you feel the force? an investigation of haptic collaboration in shared editors," in *EuroHaptics*, 2001, pp. 54–59.

[3] L. Rosenberg, "Virtual fixtures: Perceptual tools for telerobotic manipulation," in *Virtual Reality Ann. Int'l Symp.*, 1993, pp. 76–82.

[4] R. B. Gillespie, M. O'Modhrain, P. Tang, D. Zaretzky, and C. Pham, "The virtual teacher," in *ASME IMECE-DSC*, vol. 64, 1998, pp. 171–178.

[5] D. J. Reinkensmeyer, "How to retrain movement after neurologic injury: a computational rationale for incorporating robot (or therapist) assistance," in *IEEE Engineering in Medicine and Biology Society Meeting*, vol. 2, 2003, pp. 1479–1482.

[6] S. J. Harkema, "Neural plasticity after human spinal cord injury: application of locomotor training to the rehabilitation of walking," *The Neuroscientist*, vol. 7, no. 5, p. 455, 2001.

[7] C. J. Winstein, P. S. Pohl, and R. Lewthwaite, "Effects of physical guidance and knowledge of results on motor learning: support for the guidance hypothesis," *Res Q Exerc Sport*, vol. 65, no. 4, pp. 316–323, 1994.

[8] J. D. Hagman, "Presentation- and test-trial effects on acquisition and retention of distance and location." *J. Exp. Psychol.: Learning, memory, and cognition*, vol. 9, no. 2, pp. 334–345, 1983.

[9] T. R. Armstrong, "Training for the production of memorized movement patterns," in *Ann. Conf. on Manual Control*, 1970.

[10] R. A. Schmidt and R. A. Bjork, "New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training," *Psychol. Science*, vol. 3, no. 4, p. 207, 1992.

[11] J. F. Israel, D. D. Campbell, J. H. Kahn, and T. G. Hornby, "Metabolic costs and muscle activity patterns during robotic- and therapist-assisted treadmill walking in individuals with incomplete spinal cord injury," *Physical Therapy*, vol. 86, no. 11, 2006.

[12] R. Shadmehr and F. Mussa-Ivaldi, "Adaptive representation of dynamics during learning of a motor task," *J. Neuroscience*, vol. 14, no. 5 II, pp. 3208–3224, 1994.

[13] K. A. Thoroughman and R. Shadmehr, "Learning of action through adaptive combination of motor primitives," *Nature*, vol. 407, no. 6805, pp. 742–747, Oct. 2000.

[14] L. M. Crespo and D. J. Reinkensmeyer, "Effect of robotic guidance on motor learning of a timing task," in *IEEE RAS & EMBS Int'l Conf. Biomedical Robotics and Biomechatronics*, 2008, pp. 199–204.

[15] M. K. O'Malley, A. Gupta, M. Gen, and Y. Li, "Shared control in haptic systems for performance enhancement and training," *J. Dynamic Syst., Meas., and Control*, vol. 128, no. 1, pp. 75–85, Mar. 2006.

[16] L. M. Crespo and D. J. Reinkensmeyer, "Haptic guidance can enhance motor learning of a steering task," *J. Motor Behavior*, vol. 40, no. 6, pp. 545–557, 2008.

[17] T. Endo, H. Kawasaki, K. Kigaku, and T. Mouri, "Transfer method of force information using Five-Fingered haptic interface robot," in *EuroHaptics*, 2007, pp. 599–600.

[18] J. Lee and S. Choi, "Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance," *Haptics Symp.*, 2010.

[19] C. B. Zilles and J. K. Salisbury, "A constraint-based god-object method for haptic display," in *IEEE/RSJ Int'l Conf. on Intelligent Robots and Syst.*, 1995, pp. 146–151.

[20] S. Nudehi, R. Mukherjee, and M. Ghodoussi, "A shared-control approach to haptic interface design for minimally invasive telesurgical training," *IEEE Trans. Control Syst. Technol.*, vol. 13, no. 4, pp. 588–592, 2005.

[21] 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.

[22] M. O'Malley and A. Gupta, "Skill transfer in a simulated underactuated dynamic task," in *IEEE Int'l Workshop on Robot and Human Interactive Communication*, 2003, pp. 315–320.