The Task-Dependent Efficacy of Shared-Control Haptic Guidance Paradigms

Dane Powell, *Student Member*, *IEEE*, and Marcia K. O'Malley, *Member*, *IEEE*

Abstract—Shared-control haptic guidance is a common form of robot-mediated training used to teach novice subjects to perform dynamic tasks. Shared-control guidance is distinct from more traditional guidance controllers, such as virtual fixtures, in that it provides novices with real-time visual and haptic feedback from a real or virtual expert. Previous studies have shown varying levels of training efficacy using shared-control guidance paradigms; it is hypothesized that these mixed results are due to interactions between specific guidance implementations ("paradigms") 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 more complex guidance paradigms. The efficacies of four common paradigms are compared in a controlled study with 50 healthy subjects and two dynamic tasks. The results show that guidance paradigms must be matched to a task's dynamic characteristics to elicit effective training and low workload. Based on these results, we provide suggestions for the future development of improved haptic guidance paradigms.

Index Terms—Shared control, haptic rendering, haptic guidance, robot-mediated training.

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), online movement planning, and adaptation to changing task conditions. Most importantly, these tasks often have 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 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. Human-mediated training would entail an expert guiding a novice subject through a task via direct physical contact. Robot-mediated training would entail an expert sharing control of a task with a novice in a virtual environment, as shown in Fig. 1. The potential advantages of robot-mediated training and shared-control guidance are discussed in Section 2.

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 feedback to the novice has been studied comparatively little. Haptic feedback from the expert and virtual environment can enhance a novice's sense of presence and cooperation [2], [3], but the efficacy of haptic "guidance" at improving training outcomes has not been thoroughly demonstrated.

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 various extant 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. The traditionally used proxy rendering model cannot easily implement more modern and complex guidance paradigms, and thus an improved shared-control proxy model is proposed in Section 4.

We hypothesize that in order for a guidance paradigm to elicit effective training, it must be properly matched to the task at hand. To this end, we compared the efficacies of four guidance paradigms at training 50 healthy subjects to perform two dynamic tasks in a controlled study in Section 5. The results presented in Section 6 demonstrate that the efficacies of the paradigms indeed depend on the task, and the implications of these findings are discussed in Section 7.

2 Robot-Mediated Training

The defining characteristic of robot-mediated training is that guidance is administered physically to a novice subject via a haptic interface. Thus, a coach might still retain high-level control over the course of training or even participate teleoperatively, but all physical interactions with the novice are mediated by the haptic interface and related control systems (the "robot").

- *The authors are with the Department of Mechanical Engineering and Materials Science, Rice University, PO Box 1892, MS 321, Houston, TX 77251-1892. E-mail: {danep, omalleym}@rice.edu.*

Manuscript received 7 Oct. 2011; revised 29 June 2012; accepted 3 July 2012. Recommended for acceptance by A. Karniel, A. Peer, O. Donchin, F.A. Mussa-Ivaldi, and G.E. Loeb.
For information on obtaining reprints of this article, please send e-mail to: toh@computer.org, and reference IEEECS Log Number THSI-2011-10-0081. Digital Object Identifier no. 10.1109/ToH.2012.40.

1939-1412/12/$31.00 © 2012 IEEE Published by the IEEE CS, RAS, & CES

Fig. 1. A subject performing a target-hitting dynamic task with the assistance of a virtual expert.

Robot-mediated training offers many potential advantages over traditional "human-mediated" training. If enough autonomy can be given to the robot or a "virtual expert," one human expert could potentially train a large number of novices simultaneously, increasing the reach of training. The virtual environment can also be quickly changed or reset in order to facilitate training and help keep the novice's attention through long training sessions. More importantly, a robot can offer objective measures of performance much more frequently than a human expert [4]. Winstein [5] and others have shown that providing accurate and timely feedback to a novice can directly improve training outcomes, and such measures can be used by a real or virtual expert to tune or adapt other aspects of the training as the novice improves over time. For instance, Li et al. [6] and Huegel and O'Malley [7] used these measures to progressively decrease the amount of guidance provided to novices as their performances improved.

2.1 Virtual Fixtures and Shared-Control Guidance

Guidance during robot-mediated training is often provided via simple perceptual overlays such as virtual fixtures. Virtual fixtures, as proposed by Rosenberg [8], are simply perceptual overlays that passively prevent subjects from entering forbidden regions of a work environment, and are most often used to constrain a novice's motions to an optimal trajectory.

Shared-control guidance is a more recently developed form of guidance that improves upon virtual fixtures by allowing a novice to share control of a system with a real or virtual expert. O'Malley et al. [9] showed that such shared-control systems were as effective as virtual fixtures at facilitating skill transfer. Traditionally, such as in fly-by-wire aircraft control systems, conflicting control inputs by multiple agents, such as a novice and an expert, are reconciled by simply averaging the inputs, which is not necessarily the best cooperation paradigm [10]. Reed and Peshkin [11] make the following point:

> Averaging the input command is a simple strategy but not necessarily the best combination since each individual's motion will be diluted. Imagine the effect if one pilot attempts to avoid an obstacle by turning to the left while the other to the right: the average effect is straight into the obstacle.

Nudehi et al. [12] proposed a similar shared-control scheme for telesurgical training that calculated a control output based on the weighted average of the control inputs of two operators.

Other implementations share characteristics of both virtual fixtures and shared-control guidance, such as the "record-and-replay" strategy used by Gillespie et al. [13] to train novices to balance an inverted pendulum.

Such "assistive" forms of guidance 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.

2.2 Problems with Traditional Guidance

A common assumption is that physically guiding a novice through the successful completion of a task will help the novice to internalize and encode that pattern, and thus help the novice to repeat the pattern on his or her own in the future. This assumption is only weakly supported by the literature in the context of rehabilitation [14], [15], [16], and has been refuted in many cases in the context of training healthy individuals [17], [18], [19]. Schmidt and Bjork [20] showed that guidance in many sorts of training (not just in visuo-motor tasks) can actually impair learning and retention, especially if provided too frequently or in a form that is too easy to use. This discrepancy between the expected and actual results of guidance-based training has come to be known as the "guidance hypothesis."

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. [21] 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 [22] 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 the loop as it would during practice.

Another problem with assistive guidance is that novices make fewer errors than they would during practice because they are passive and constrained to an optimal trajectory. Thoroughman and Shadmehr [23] 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, we hypothesize that a significant problem with traditional 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

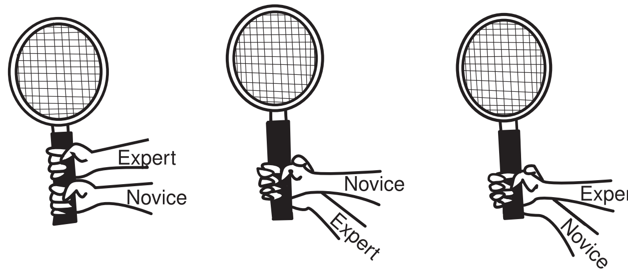

Fig. 2. Gillespie et al.'s Virtual Teacher paradigms [13]. From left to right: indirect-contact, double-contact, and single-contact paradigms.

and guidance forces. In the same way that simply averaging the command inputs of a novice and expert may have an undesirable effect on a system, averaging the task dynamics and guidance forces may have an undesirable effect on a novice's training. 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 [24], 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."

2.3 Biomimetic Forms of Guidance

Gillespie et al. [13] proposed the use of a virtual teacher, a more active form of guidance than virtual fixtures, which instructs novice subjects to perform dynamic tasks by giving them shared control of a task with a virtual expert. 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 Fig. 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. [13] 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.

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 using a consolidated and unified set of descriptors as described in this section. 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).

This taxonomy (first proposed by Powell and O'Malley [1]) classifies paradigms along three dimensions or factors. The principal factor differentiating guidance paradigms is whether they assist or resist the novice subject in completing the task. Guidance schemes can thus be classified as either "assistive" or "resistive." The second factor is how paradigms reconcile the copresentation of task and guidance forces. As mentioned in Section 1, task and guidance forces should be interpreted by the novice in fundamentally different ways. If the novice cannot clearly distinguish between the two, the guidance forces will alter the perceived dynamics of the task and potentially impair training. Most existing guidance schemes confound task and guidance forces in just such a way by combining them using a simple weighted average function so that both forces can be displayed simultaneously via a single haptic device in manner that we will refer to as "gross" guidance. Finally, many guidance schemes will adjust the relative weights (gains) of these forces over time in response to a novice's performance improvement. We refer to such schemes as "progressive."

In the following sections, we pick five representative guidance paradigms from this taxonomy and discuss their existing implementations as described in literature.

3.1 Gross Assistance (GA)

Classic virtual fixtures are the archetypal example of gross assistance. 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 toward 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 [16]. Reinkensmeyer [14] showed in simulation that "continual guidance" (GA) is "never beneficial compared to no assistance." Marchal-Crespo and Reinkensmeyer [25] 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. "Triggered" assistance is a type of GA that requires the novice to exert a certain amount of control effort before assistance is provided, and has not been shown to be conclusively better than standard GA. O'Malley et al. [26] implemented a force-based triggered mode on the MIME/RiceWrist exoskeleton, while Kahn et al. [27] implemented a displacement-based triggered mode on the ARM Guide, but neither showed any significant improvement over practice for the rehabilitation of stroke patients.

Generally speaking, most of the assistive paradigms discussed in Section 2 that can be classified as GA were shown to be ineffective compared to practice without guidance. This negative outcome is partially predicted and explained by the "guidance hypothesis" proposed by Salmoni et al. [28], which states that subjects will tend to become reliant on guidance when it is present in order to improve performance instead of relying on "other cues in the task that are important for motor learning." Li et al. [29] also found evidence of subjects becoming reliant on GA guidance forces. Furthermore, we predict that even for tasks where guidance forces do not dominate the inherent task dynamics and subjects are not passive, GA may still impair training as described in Section 2.

One possible exception to the generally negative efficacy of GA is for tasks that are extraordinarily difficult and for novices in the very early stages of training for a new task. Marchal-Crespo and Reinkensmeyer [25] showed that there was a significant improvement of the GA groups over the practice groups in the very first stages of training, but that this improvement quickly diminished and became insignificant as training continued.

Progressive gross assistance (PGA) capitalizes on the early-stage benefit of GA by systematically decreasing the guidance gains over time as a novice's performance improves, allowing the novice to make more errors in later stages of training and further refine his or her motor control. Guidance may be decreased either on a predetermined schedule or "adaptively" in response to a subject's performance. Many of the same studies in Section 2 showing that GA was ineffectual also showed that PGA was superior to both GA and practice [16].

However, PGA has some potential downfalls. First, PGA requires gain-reduction algorithms that may depend on accurate and objective performance metrics. Choosing the correct algorithm and performance metrics is highly task-dependent and potentially difficult. As with traditional GA, PGA confounds guidance and task forces during the majority of training, and may in fact exacerbate impairment by subtly changing guidance gains (and thus the task dynamics as well) over time. For these reasons, we chose not to evaluate PGA in this study.

3.2 Temporally Separated Assistance (TSA)

Temporally separated assistance separates guidance and task forces temporally, displaying each type alternately in quick succession via a single haptic device. In other words, novices are "nudged" toward the virtual expert by brief pulses of guidance provided on the order of 1 Hz. In this way, the guidance exerts "cognitive dominance" over the novice, while allowing the novice to retain "physical dominance," commit errors, and actively generate movement plans in order to better learn the task dynamics. With this advantage, we hypothesize that TSA can achieve the same level of performance as PGA without being subject to the complexities of adaptive algorithms. Additionally, compared to progressive paradigms that provide all of the guidance during training "up front," TSA provides guidance consistently and predictably throughout training, hopefully improving training outcomes.

In a pilot study, Endo et al. [30] 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. Ahn and Hogan [31] also implemented TSA in order to study entrainment of human gait, and found that the presence of properly designed TSA could encourage subjects to adopt certain gait patterns. However, they did not study TSA in the context of training.

3.3 Spatially Separated Assistance (SSA)

Whereas TSA separates the presentation of task and guidance forces temporally in order to present them via a single haptic channel, spatially separated assistance 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. [13], 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. [13] could not conclusively show that SSA was superior to practice.

Similarly, Wulf et al. [32] showed that a weak form of SSA was superior to practice without physical guidance at training novices to perform a simulated skiing task. This might be considered a "weak" form of SSA because haptic feedback was provided via actual mechanical fixtures rather than electromechanical systems and a virtual expert. However, this guidance paradigm still qualifies as SSA because guidance was provided via a spatially distinct channel (i.e., the poles) from the primary interface with the simulator (i.e., the skis).

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. Numerous studies indicate that taking advantage of bimanual (mirror) symmetry can improve rehabilitation from hemiparesis following stroke [33], [34]. Additionally, studies have shown that in healthy individuals, there is a transference of skills between bimanual and unimanual tasks [35]. Finally, Tcheang et al. [36] showed that forces applied to one arm will not interfere learning of force fields by the other arm. These studies support our implementation of SSA as described in Section 5.4.

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 overtraining: 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.

Constant (Coulomb) or velocity-dependent (viscous) forces opposing the direction of movement have been shown to improve training outcomes particularly in the field of rehabilitation. For instance, Lambercy et al. [37] designed a haptic knob offering varying levels of resistive force in order to help stroke patients regain grasp strength and coordination. A metareview by Morris et al. [38] showed that resistance training (though not necessarily robot-mediated) can help reduce musculoskeletal impairment in stroke patients.

Error augmentation has also been shown to improve training by taking advantage of the CNS' error-driven learning process. Emken and Reinkensmeyer [39] showed that amplifying the task dynamics and in turn producing larger movement errors improved the adaptation of healthy novices to a viscous force-field. In rehabilitation, Patton et al. [40] showed that force-fields that amplified movement errors made by stroke patients in a reaching task improved training outcomes over practice.

Finally, Lee and Choi [41] showed that training in the presence of random noise-based disturbance was superior to PGA 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

A number of factors make stable rendering of the interaction between a haptic interface and virtual environment a nontrivial task. Foremost among these are the physical limitations of even the most modern haptic devices, which tend to be relatively compliant compared to the virtual objects that they interact with. In order to maintain a one-to-one relationship between the position of the haptic device in real space and in the virtual environment, the device would have to penetrate unrealistically far into the virtual object. Thus, direct calculation of interaction forces based on a physics model is generally not possible, as the forces would tend to saturate quickly enough to lead to explosive instability, and some other general haptic rendering algorithm is required.

Zilles and Salisbury [42] 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 human 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.

If a perceptual overlay or virtual expert is added to the environment, one can imagine that two qualitatively different types of forces exist 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. This distinction is important because guidance forces should be used to shape the novice's actions, whereas task forces should be incorporated into the novice'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 novice. This precludes the use of the more advanced separation guidance paradigms described in Section 3.

Step 1: Compute proxy location
x_p = f(x_e, x_n, α)
α = f(k_e, k_n, b_e, b_n)
Step 2: Compute guidance force
F_G = f(x_p, x_n)
Step 3: Compute task force
F_T = f(x_p, x_m, k_m, b_m)
Step 4: Compute total force displayed to novice
F_{out} = f(F_G, F_T)

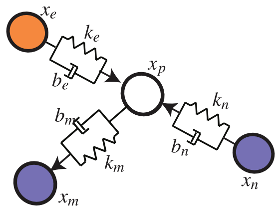

Fig. 3. Shared-control proxy model algorithm. Variables x_e, x_n, x_p, and x_m represent the positions of the expert, novice subject, proxy, and mass.

Our shared-control proxy model (first proposed by Powell and O'Malley [1]) 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 Fig. 3, where arrows indicate the general directions of force transfer.

Although modeled as a system of springs and dampers, it is easiest to understand the operation of the shared-control proxy model by thinking about it algorithmically:

1. Compute the proxy position based on the positions of the novice and expert. The proxy represents the average input of the novice and expert, weighted according to a control authority α (as proposed by Nudehi et al. [12]); thus, the proxy *always* lies on a line connecting the novice and expert.
2. Compute the guidance force based on the displacement between the proxy and the novice; the larger this displacement, the farther the novice is straying from the expert's position.
3. Compute the task force based on the displacement between the proxy and the mass, which interacts with the environment. Thus, as the proxy penetrates a virtual surface, this displacement and task force will grow.
4. Compute the output force to the novice based on a combination of the guidance and task forces. For this study, the output was calculated as shown in Table 1.

Kucukyilmaz et al. [43] and Oguz et al. [44] proposed a similar rendering method designed to facilitate role-exchange in shared-control systems. Their role-exchange

TABLE 1
Force Outputs during Training

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

The task force F_T and guidance force F_G are calculated using the shared-control proxy model (Section 4). F_{PN} is given by a Perlin noise function. $t_0 = 100$ ms *and* $t_1 = 500$ ms.

model provides only task forces to a user, while our shared-control proxy model provides task forces and guidance forces separately, allowing them to be modulated using either SSA or TSA. This is made possible through the use of variable-ratio virtual couplings, as opposed to the simple spring-damper couplings in the role-exchange model.

5 Methods

We evaluated the effectiveness of four prototypical guidance schemes at training 50 novice subjects to perform two dynamic tasks in a controlled study. Subjects controlled the position of an on-screen pointer using a 2-DOF haptic joystick (Immersion, Inc.'s, IE2000), where the subject's ulnar/radial deviation and pronation/supination are mapped to cursor position. The maximum force output of the joystick was 5 N. Subjects trained with the assistance of a virtual expert using the shared-control proxy model shown in Fig. 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 1,000 Hz, the visual display was rendered by OpenGL at 60 Hz, and experimental data was recorded at 100 Hz.

5.1 Experimental Design

Subjects performed the two tasks on two consecutive days, with a single session and type of task per day (the order of task presentation was balanced between groups). Each 1-hour session consisted of 106 trials grouped into a number of blocks, as shown in Fig. 4. Subjects were allowed a 1-minute familiarization trial with an easier version of the task prior to each session, as well as a 5-minute break midway through the session. This experimental design prevented subject fatigue while minimizing scheduling burdens.

The bulk of each session consisted of 20-second-long "evaluation" and "training" trials. 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. In training trials, subjects shared control of the system with a virtual expert under one of the guidance paradigms described in Section 5.4. The virtual expert followed an optimal pre-computed trajectory that was identical between subjects. Subjects were instructed to track the expert as closely as possible during training and replicate its behavior during evaluation.

Interspersing evaluation and training trials in this manner made it possible to record how subjects' performances improved over time, as opposed to simpler experimental designs with only pretraining and posttraining evaluations.

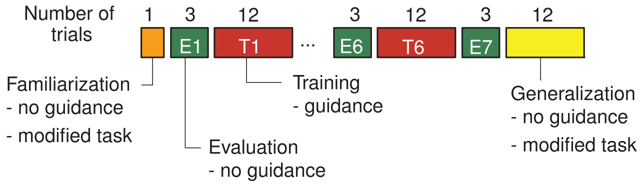

Fig. 4. Illustration of session structure for each task.

This design also made it possible to ensure that subjects achieved complete training, by observing that their performances plateaued at the end of each session.

Finally, generalization trials were presented at the end of each session in order to test the robustness of acquired motor skills to changing task dynamics. Subjects also reported their perceived task workload by completing the NASA TLX questionnaire developed by Hart [45]. This questionnaire allows subjects to rate their perceived workload on six different subscales: 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.

5.2 Subjects

A total of 50 subjects enrolled in the primary study, and were divided evenly between five experimental groups: visual-only 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. Subjects were instructed only on their specific guidance condition; they were not made aware of guidance paradigms besides their own, or of whether they were part of the control group.

5.3 Tasks

5.3.1 *Target-Hitting Task*

The target-hitting task used in these experiments was based on a task originally used by O'Malley and Gupta [46] and O'Malley et al. [9]. The novice's proxy was connected to a 5 kg mass by a spring with stiffness $k = 100$ N/m and damping $b = 3$ Ns/m, as shown in Fig. 5. Two targets were positioned equidistant from the center of the screen and at a 45 degree 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 activate. Each task trial was 20 seconds long, and the goal during evaluation trials was to hit as many targets as possible. 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 hit-count corresponding to the optimum trajectory (28 hits).

During training, subjects shared control of this system with a virtual expert via the shared-control proxy model described in Section 4. The virtual expert was represented on-screen by an orange pointer that tracked the optimal

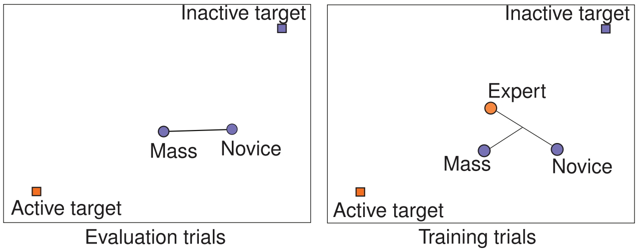

Fig. 5. Visual display for target-hitting task. Subjects control a pointer that is linked to a mass-spring-damper system and try to guide the mass to alternating targets. In evaluation trials, a simple spring and damper link the mass and the novice. In training trials, the shared-control proxy model links the novice, expert, proxy, and mass.

trajectory, a sinusoidal movement along the straight-line axis between the targets with a frequency of 0.71 Hz. Fig. 1 illustrates a subject performing this task.

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. Gift cards were awarded to the subjects that best achieved these goals.

For generalization trials, the task parameters were changed so that the mass $m = 2\ \mathrm{kg}$, stiffness $k = 80\ \mathrm{N/m}$, and damping $b = 1\ \mathrm{Ns/m}$.

5.3.2 Path-Following Task

The path-following task was similar to a traditional pursuit-rotor task in that it required novices to track a virtual expert around the outline of a simple shape at a constant speed. The task is based loosely on that proposed by Lee and Choi [41]. In evaluation and training trials, subjects traced the outline of one of the three shapes shown in Fig. 6 (a circle, square, or X). In any given block, the shapes would be presented in equal number but a random order. In the familiarization and generalization blocks, subjects were shown triangles and lemniscates (respectively). In all cases, the goal was to trace the expert as closely as possible, and thus performance was defined as cumulative deviation from the expert's position (in centimeters) over the course of each trial. Gift cards were awarded to the subjects that achieved the lowest deviation.

5.4 Haptic Guidance Paradigms

Visual guidance was always provided during training trials. Haptic guidance was provided using one of the guidance paradigms described below and in Table 1.

5.4.1 Visual-Only 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 its position.

5.4.2 Gross Assistance

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

Similar to GA, task and guidance forces were combined by summation and presented via a single joystick. However,

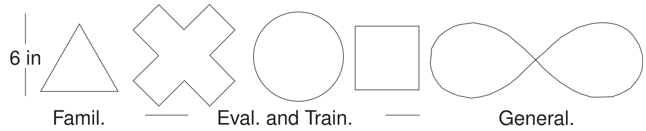

Fig. 6. Shapes used in the path-following task.

guidance forces were modulated at a frequency of 2 Hz, so that subjects experienced 100 ms pulses of guidance followed by 400 ms of completely attenuated (zero) guidance each period. This is the optimal frequency and ratio as experimentally derived by Endo et al. [30]. 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.

5.4.4 Spatially Separated Assistance

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 nondominant 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 of the task and not rely on the guidance to do any "heavy lifting."

5.4.5 Gross Resistance

Task forces were combined with a randomly generated disturbance force in the manner described by Lee and Choi [41]. A Perlin noise function with a nominal range of -1.2 to 1.2 N was randomly generated for each joystick axis using the open-source *libnoise* library. At each time step, the guidance force generated by these functions was summed with the task force to produce the net force displayed to the joystick.

6 Results

Performance for almost all subjects approached an asymptotic "saturation" level by the end of each session, indicating that complete training was achieved for both tasks.

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 the respective cell mean. Vertical lines on column graphs indicate standard error.

Results are reported for mixed ANOVA omnibus and interaction tests, as well as for family-wise error-corrected (Tukey-Kramer (TK) adjusted) multiple comparisons. Horizontal lines between groups indicate TK-adjusted significance at $\alpha = .05$. The presence of an interaction effect between group and trial number indicates that performance between groups differs depending on trial number (level of training).

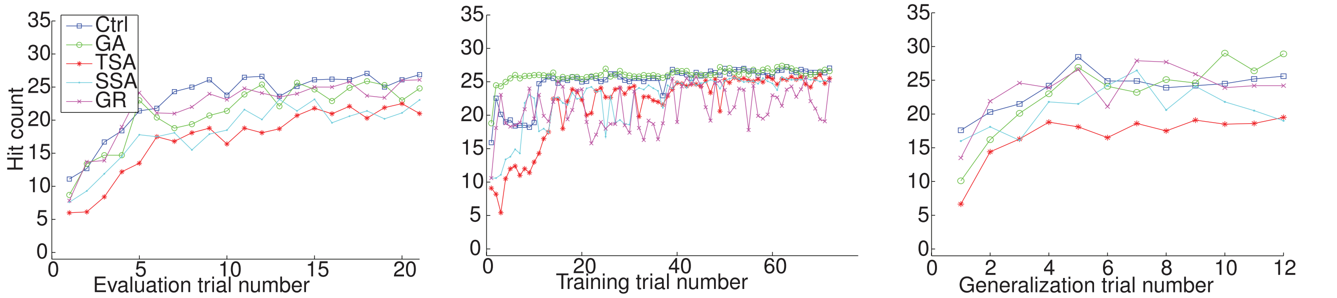

(a) (b) (c)

Fig. 7. Subject hit counts, averaged by guidance group, for target-hitting trials.

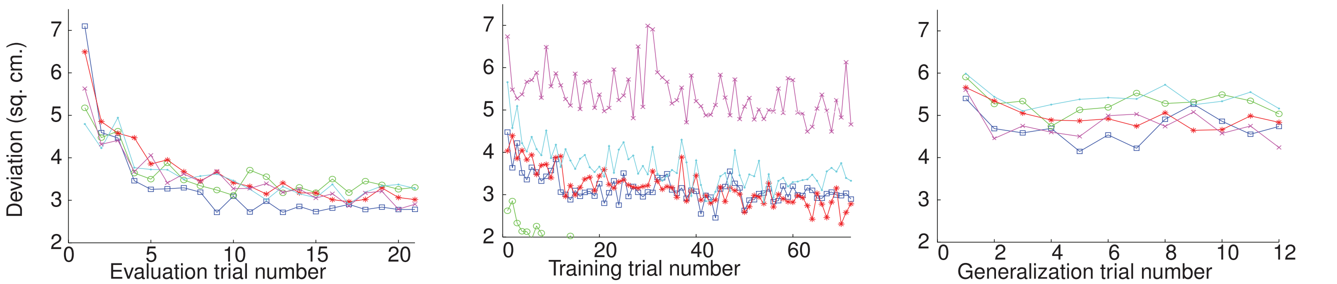

(a) (b) (c)

Fig. 8. Subject deviations, averaged by guidance group, for path-following trials.

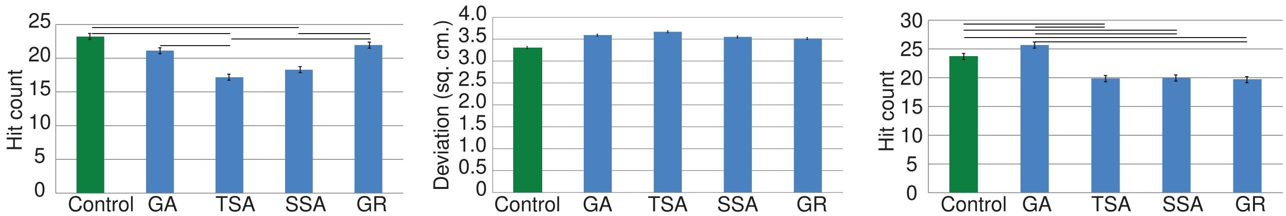

(a) Target-hitting evaluation trials. Fixed effect of group: $F(4,113) = 9.46$, $p < .001$.

(b) Path-following evaluation trials. Fixed effect of group: $F(4,88) = 1.42$, $p = .219$. Interaction effect of group and trial: $F(80,852) = 1.43$, $p = .01$.

(c) Target-hitting training trials. Fixed effect of group: $F(4,104) = 4.60$, $p = .002$. Interaction effect of group and trial: $F(284,3061) = 2.31$, $p < .001$.

Fig. 9. Results of mixed ANOVA fixed effect and interaction tests.

6.1 Mixed ANOVA on Evaluation Trials

The performance of each group during evaluation trials is shown in Figs. 7a and 8a. Mixed ANOVAs were performed on evaluation trial outcomes with hit counts (for target-hitting) or deviation (path-following) as the dependent variable, guidance paradigm ("group") as a between-subjects factor, and trial number ("trial") as the within-subjects factor (repeated measure). Outcomes for the omnibus ANOVA and for pairwise multiple comparisons, corrected using a TK adjustment, are shown for target-hitting and path-following in Figs. 9a and 9b, respectively.

All groups exhibited a consistent learning trend. Multiple comparisons based on the mixed ANOVA showed that for target-hitting 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.

6.2 Mixed ANOVA on Training Trials

Although the primary goal of this research is to improve robot-mediated training methodologies, the guidance paradigms being tested could also be used for online correction of human inputs in the midst of task execution, such as when the autopilot or stick shaker mechanism in an aircraft shares control with the human pilot. Thus, it is useful to know how each guidance paradigm affects task performance while the guidance is actually active (during training), as shown in Figs. 7b and 8b. Fig. 9c shows that TSA, SSA, and GR all performed significantly worse than the control group and GA.

6.3 Mixed ANOVA on Generalization Trials

Also of interest is how motor learning effects transfer to similar tasks with slightly modified task dynamics. The results from the generalization trials reflect those of the evaluation trials, in that the TSA group performed poorly in target-hitting (Fig. 7c), while the GA and SSA groups performed poorly in path-following (Fig. 8c). The omnibus ANOVA was not statistically significant, however.

6.4 Workloads

Workloads were recorded on six subscales, as described in Section 5.1; these workloads are shown in Figs. 10 and 11, normalized about the control group for easy comparison between groups and tasks. Note, however, that it is *not* valid to compare the workload values between subscales, since each subscale contributes to the overall workload differently. For instance, even though the highest workload

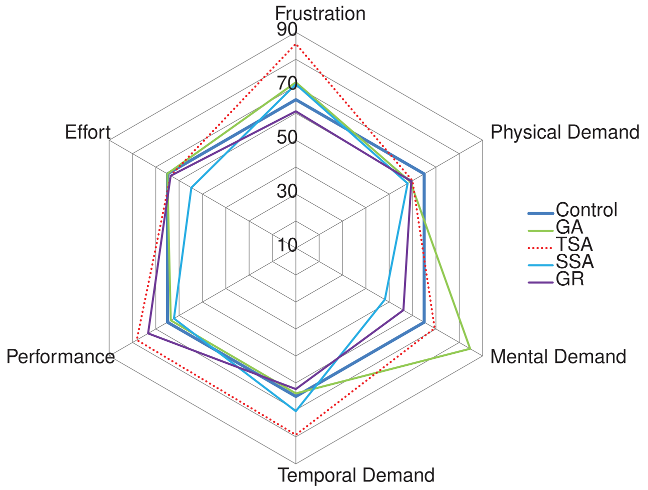

Fig. 10. Target-hitting task workloads, normalized to the control group.

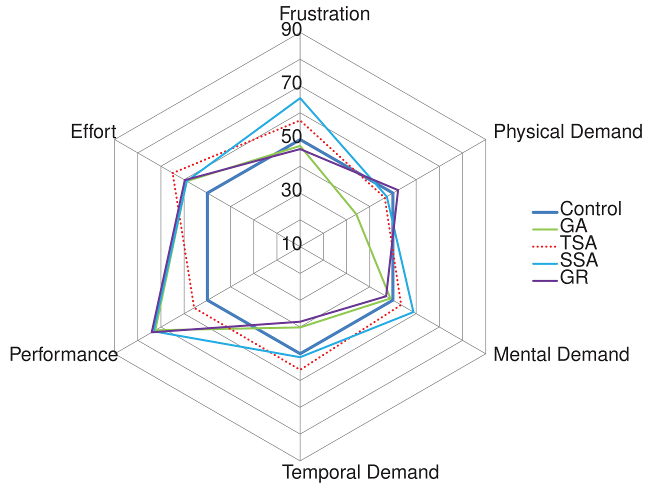

Fig. 11. Path-following task workloads, normalized to the control group.

in the target-hitting task was on the frustration subscale, subjects indicated that their frustration contributed very little to their overall sense of workload in the task.

In the target-hitting task, none of the guidance schemes significantly increased effort or physical demand compared to the control group. SSA produced notably less mental demand and effort, while TSA produced notably higher frustration and GA notably higher mental demand. Post-study interviews with subjects in the GA group indicate that they found the assistance very confusing, as described in Section 7. GR actually decreased frustration, physical and mental demand, and temporal demand compared to the control group.

By contrast, in the path-following task, all of the guidance schemes actually increased effort and decreased subjects' perceived performance. TSA and SSA also led to higher frustration. However, GA and GR led to notably less temporal and mental demand, and GA also reduced the physical demand.

The path-following task had a significantly lower workload than the target-hitting task in almost every respect. Permutation tests indicate that many subscale pairwise comparisons were statistically significant. Permutation tests were also performed on overall workload scores, but no pairwise comparisons survived a TK adjustment.

7 Discussion

7.1 Efficacies of Guidance Paradigms

The target-hitting and path-following tasks each have unique dynamic characteristics. The target-hitting task is temporally demanding, has high levels of inherent task forces, and requires novices to optimize their excitation frequency of the system. By contrast, the path-following task is slower, lacks inherent task forces, and places more emphasis on precise position control. By comparing the efficacies of the guidance paradigms between these two tasks, we can draw conclusions about the interplay between task characteristics, choice of guidance paradigm, and training efficacy.

Our results corroborate the guidance hypothesis [20], which indicates that challenge is necessary for the learning process, and show that this effect is exacerbated in a task dominated by guidance forces. In the path-following task, which had no inherent task forces, subjects trained with GA reported significantly lower physical demand but performed poorly in evaluation trials. These data indicate that novices have greater passivity in tasks dominated by guidance forces, leading to poor learning as the guidance hypothesis predicts.

Our results also confirm our hypothesis that even in tasks with significant inherent forces, GA leads to confusion by making task and guidance forces indistinguishable. In contrast to the path-following results, GA subjects reported extremely high mental demand in the target-hitting task and reported that the forces during training were "confusing" and "difficult to interpret." However, in evaluation trials, subjects trained with GA appeared to perform better than those trained with TSA and not significantly worse than the control group.

The poor performance of the separation paradigms is surprising, given that they were specifically designed to discourage dependency without interfering with inherent task dynamics. Many subjects reported that they found the constant "nudging" of TSA to be frustrating, which was reflected in the workload results. Additionally, it is possible that the poor performance of TSA in the target-hitting task was 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 optimum hit count in the task, following the expert is not a necessary condition for achieving the optimum hit count. It is possible to follow the expert at a phase lag and still achieve the optimum hit count—in fact, guidance forces and task forces are equal and opposite when the novice is out of phase with the expert by a certain amount, leading to the confusion in the GA group discussed above.

This corroborates the work of Ahn and Hogan [47], who found that TSA was only effective if the frequency and phase shift of the guidance matched critical values unique to the task. Specifically, they found that entrainment, or adoption of a new gait pattern, only occurred when the perturbation frequency was within 7 percent of the natural cadence, and when perturbations were provided during the heel-strike portion of the gait cycle.

Our findings regarding GR somewhat support those of Lee and Choi [41]. On a very similar path-following task, they found that training with GR led to significantly better performance than PGA. In this study, the results suggest that GR outperformed GA in path-following, especially in the generalization trials (Fig. 8c), although these results were not statistically significant. In terms of workload, GR led to higher effort and physical demand in path-following, but smaller workloads in nearly every respect in target-hitting.

Overall, our results confirm the guidance hypothesis as well as our own hypothesis regarding the inadequacies of the GA paradigm, and show that the training efficacy of guidance is highly task dependent. Our results also suggest that factors besides just challenge level and passivity can have a substantial effect on training efficacy, given the poor performance of the TSA and SSA paradigms.

7.2 Effects during Training

Considering just the effect of the guidance paradigms on performance during training, it is worthwhile to note that the control and GA groups outperformed most other groups, and results suggest that the GA group outperformed the control group. This suggests that if guidance is being used to assist an operator in the real-time execution of tasks, for instance to prevent the operator from entering dangerous or forbidden regions of the workspace, then GA is the guidance method of choice. Additionally, the positive effect of interaction between hit count and trial number indicates that the performance of subjects receiving gross guidance (GA or GR) quickly plateaus, while subjects receiving separated guidance (TSA or SSA) steadily improve over the course of the study. This suggests that the more complex separated guidance paradigms are less intuitive than the simpler gross paradigms, and perhaps require more instruction on proper utilization.

7.3 Generalizability

We believe that these results should generalize to most types of dynamic tasks, based on the nature of the tasks in this study and the results of the generalization trials. The tasks are characteristically different in nature in multiple categories: target-hitting is a semidiscrete rhythmic task with inherent haptic feedback, whereas the path-following task is continuous and lacks haptic task feedback. Within each type of task, results in the generalization trials were not significantly different from those in evaluation trials.

7.4 Future Directions

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 in the target-hitting task or augmenting the perceived error in the path-following task might both be effective ways of enhancing training.

There are several take-away lessons to be applied to future studies. First, some subjects seemed to be "natural experts," performing well in the initial evaluation and improving little over the course of the study, while others performed very poorly at entrance. Thus, it is important to balance the distribution of subjects between groups based, for instance, on performance during the familiarization task.

Second, it is important that the guidance paradigms be applied to tasks of sufficient difficulty to elicit long-duration improvement. If the task is too easy, any effects of training may be overshadowed by the effects of subjects reaching the performance ceiling too soon. Similarly, when designing TSA, it is important to properly match the frequency and phase of the guidance to the task, as described in Section 7.

Third, in regards to workloads, many subjects reported that they did not understand the pairwise comparisons between subscales, and did not have a reference on which to base their reported workload on each subscale. In the future, more intuitive methods of assessing workload, such as a qualitative scale (i.e., "could perform the task in my sleep" through "task is physically impossible"), might improve power. For instance, Tsang and Velazquez [48] developed the Workload Profile scale, which Rubio et al. [49] found to have higher sensitivity and diagnostic power than NASA TLX. Variances due to self-reporting could also be avoided by using secondary tasks to assess workload.

8 Conclusions

We have shown that properly matching a guidance paradigm to a task's dynamic characteristics is critical for achieving high efficacy and low workload, and that many types of guidance can actually impair training as compared to practice. Previous studies have shown that subjects can become dependent on assistive guidance, and we have shown that this is exacerbated in a task without inherent dynamic forces. We have also shown that assistive guidance may impair training in more dynamic tasks by altering the task dynamics.

The results of this study also indicate that challenge is indeed essential to the learning process, as predicated by the guidance hypothesis, but that increasing the challenge level via resistive guidance does not necessarily improve training. However, we have shown that GR, when properly matched to a task, does not have a negative effect on training, and may in fact have the beneficial effect of lowering workload.

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 based on this model. Although these paradigms were not shown to be superior in the context of training, it is hypothesized that they would be very advantageous in the context of robot-mediated real-time task execution, and could improve shared-control human-machine interfaces ranging from autopilots to robotic surgical systems.

Acknowledgments

Supported by US National Science Foundation (NSF) Grant IIS-0448341. Portions of this work have been previously published [1]. The authors wish to thank Rebecca Lam for her illustrations.

References

[1] D. Powell and M.K. O'Malley, "Efficacy of Shared-Control Guidance Paradigms for Robot-Mediated Training," *Proc. IEEE World Haptics Conf.,* 2011.

[2] C. Ho et al., "An Experiment on the Influence of Haptic Communication on the Sense of Being Together," *Proc. BT Presence Workshop,* 1998.

[3] I. Oakley, S. Brewster, and P. Gray, "Can you Feel the Force? An Investigation of Haptic Collaboration in Shared Editors," *Proc. EuroHaptics Conf.,* pp. 54-59, 2001.

[4] O. Celik et al., "Comparison of Robotic and Clinical Motor Function Improvement Measures for Sub-Acute Stroke Patients," *Proc. IEEE Int'l Conf' Robotics and Automation (ICRA),* 2008.

[5] C.J. Winstein, "Knowledge of Results and Motor Learning Implications for Physical Therapy," *Physical Therapy,* vol. 71, no. 2, pp. 140-149, 1991.

[6] Y. Li et al., "Progressive Shared Control for Training in Virtual Environments," *Proc. Third Joint EuroHaptics Conf. and Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC '09),* 2009.

[7] J. Huegel and M. O'Malley, "Visual Versus Haptic Progressive Guidance for Training in a Virtual Environment," *Proc. Third Joint Eurohaptics Conf. and Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC '09),* 2009.

[8] L. Rosenberg, "Virtual Fixtures: Perceptual Tools for Telerobotic Manipulation," *Proc. Virtual Reality Ann. Int'l Symp.,* pp. 76-82, 1993.

[9] M.K. O'Malley et al., "Shared Control in Haptic Systems for Performance Enhancement and Training," *J. Dynamic Systems, Measurement, and Control,* vol. 128, no. 1, pp. 75-85, Mar. 2006.

[10] L.G. Summers et al., "Fly-by-Wire Sidestick Controller Evaluation," SAE Technical Paper Series, NASA, 1987.

[11] K. Reed and M. Peshkin, "Physical Collaboration of Human-Human and Human-Robot Teams," *IEEE Trans. Haptics,* vol. 1, no. 2, pp. 108-120, July-Dec. 2008.

[12] S. Nudehi, R. Mukherjee, and M. Ghodoussi, "A Shared-Control Approach to Haptic Interface Design for Minimally Invasive Telesurgical Training," *IEEE Trans. Control Systems Technology,* vol. 13, no. 4, pp. 588-592, July 2005.

[13] R.B. Gillespie et al., "The Virtual Teacher," *Proc. ASME Int'l Mechanical Eng. Congress and Exposition (IMECE-DSC),* vol. 64, pp. 171-178, 1998.

[14] D.J. Reinkensmeyer, "How to Retrain Movement After Neurologic Injury: A Computational Rationale for Incorporating Robot (or Therapist) Assistance," *Proc. IEEE 25th Ann. Int'l Conf. Eng. in Medicine and Biology Soc. Meeting,* vol. 2, pp. 1479-1482, 2003.

[15] 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, pp. 455-468, 2001.

[16] L. Marchal-Crespo and D. Reinkensmeyer, "Review of Control Strategies for Robotic Movement Training After Neurologic Injury," *J. NeuroEng. and Rehabilitation,* vol. 6, no. 1, p. 20, 2009.

[17] 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," *Research Quarterly for Exercise and Sport,* vol. 65, no. 4, pp. 316-323, 1994.

[18] J.D. Hagman, "Presentation- and Test-Trial Effects on Acquisition and Retention of Distance and Location," *J. Experimental Psychology: Learning, Memory, and Cognition,* vol. 9, no. 2, pp. 334-345, 1983.

[19] T.R. Armstrong, "Training for the Production of Memorized Movement Patterns," *Proc. Ann. Conf. Manual Control,* 1970.

[20] R.A. Schmidt and R.A. Bjork, "New Conceptualizations of Practice: Common Principles in Three Paradigms Suggest New Concepts for Training," *Psychological Science,* vol. 3, no. 4, pp. 207-217, 1992.

[21] J.F. Israel et al., "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, pp. 1466-1478, 2006.

[22] 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.

[23] 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.

[24] L.M. Crespo and D.J. Reinkensmeyer, "Effect of Robotic Guidance on Motor Learning of a Timing Task," *Proc. IEEE RAS and EMBS Int'l Conf. Biomedical Robotics and Biomechatronics,* pp. 199-204, 2008.

[25] L. Marchal-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.

[26] M.K. O'Malley et al., "The RiceWrist: A Distal Upper Extremity Rehabilitation Robot for Stroke Therapy," *Proc. ASME Int'l Mechanical Eng. Congress and Exposition,* 2006.

[27] L.E. Kahn et al., "Robot-Assisted Reaching Exercise Promotes Arm Movement Recovery in Chronic Hemiparetic Stroke: A Randomized Controlled Pilot Study," *J. NeuroEng. and Rehabilitation,* vol. 3, pp. 1-13, 2006.

[28] A.W. Salmoni, R.A. Schmidt, and C.B. Walter, "Knowledge of Results and Motor Learning: A Review and Critical Reappraisal," *Psychological Bull.,* vol. 95, no. 3, pp. 355-386, 1984.

[29] 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. Applied Perception,* vol. 6, no. 1, pp. 1-21, 2009.

[30] T. Endo et al., "Transfer Method of Force Information Using Five-Fingered Haptic Interface Robot," *Proc. Second Joint Eurohaptics Conf. and Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC '07),* pp. 599-600, 2007.

[31] J. Ahn and N. Hogan, "The Basin of Entrainment of Human Gait under Mechanical Perturbation," *Proc. Dynamic Systems and Control Conf. (DSCC '08),* 2008.

[32] G. Wulf, C.H. Shea, and C.A. Whitacre, "Physical-Guidance Benefits in Learning a Complex Motor Skill," *J. Motor Behavior,* vol. 30, no. 4, pp. 367-380, Dec. 1998.

[33] J. Whitall et al., "Repetitive Bilateral Arm Training with Rhythmic Auditory Cueing Improves Motor Function in Chronic Hemiparetic Stroke," *Stroke,* vol. 31, no. 10, pp. 2390-2395, Oct. 2000.

[34] H.G. Malabet, R.A. Robles, and K.B. Reed, "Symmetric Motions for Bimanual Rehabilitation," *Proc. IEEE/RSJ Int'l Conf. Intelligent Robots and Systems,* pp. 5133-5138, Oct. 2010.

[35] D. Nozaki, I. Kurtzer, and S.H. Scott, "Limited Transfer of Learning between Unimanual and Bimanual Skills within the Same Limb," *Nature Neuroscience,* vol. 9, no. 11, pp. 1364-1366, Oct. 2006.

[36] L. Tcheang et al., "Simultaneous Bimanual Dynamics are Learned Without Interference," *Experimental Brain Research,* vol. 183, no. 1, pp. 17-25, Oct. 2007.

[37] O. Lambercy et al., "A Haptic Knob for Rehabilitation of Hand Function," *IEEE Trans. Neural Systems Rehabilitation Eng.,* vol. 15, no. 3, pp. 356-366, Sept. 2007.

[38] S.L. Morris, K.J. Dodd, and M.E. Morris, "Outcomes of Progressive Resistance Strength Training Following Stroke: A Systematic Review," *Clinical Rehabilitation,* vol. 18, no. 1, pp. 27-39, Jan. 2004.

[39] J.L. Emken and D.J. Reinkensmeyer, "Robot-Enhanced Motor Learning: Accelerating Internal Model Formation During Locomotion by Transient Dynamic Amplification," *IEEE Trans. Neural Systems Rehabilitation Eng.,* vol. 13, no. 1, pp. 33-39, Mar. 2005.

[40] J.L. Patton, M. Kovic, and F.A. Mussa-Ivaldi, "Custom-Designed Haptic Training for Restoring Reaching Ability to Individuals with Stroke," *J. Rehabilitation Research and Development,* vol. 43, pp. 643-656, 2006.

[41] J. Lee and S. Choi, "Effects of Haptic Guidance and Disturbance on Motor Learning: Potential Advantage of Haptic Disturbance," *Proc. Haptics Symp.,* 2010.

[42] C.B. Zilles and J.K. Salisbury, "A Constraint-Based God-Object Method for Haptic Display," *Proc. IEEE/RSJ Int'l Conf. Intelligent Robots and Systems,* pp. 146-151, 1995.
[43] A. Kucukyilmaz, T.M. Sezgin, and C. Basdogan, "Conveying Intentions through Haptics in Human-Computer Collaboration," *Proc. IEEE World Haptics Conf. (WHC),* pp. 421-426, http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5945523, 2011.
[44] S.O. Oguz et al., "Haptic Negotiation and Role Exchange for Collaboration in Virtual Environments," *Proc. IEEE Haptics Symp.,* pp. 371-378, http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5444628, 2010.
[45] S.G. Hart, "Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research," technical report, NASA Ames Research Center, 1988.
[46] M. O'Malley and A. Gupta, "Skill Transfer in a Simulated Underactuated Dynamic Task," *Proc. IEEE Int'l Workshop Robot and Human Interactive Comm.,* pp. 315-320, 2003.
[47] J. Ahn and N. Hogan, "Feasibility of Dynamic Entrainment with Ankle Mechanical Perturbation to Treat Locomotor Deficit," *Proc. IEEE Ann. Int'l Conf. Eng. in Medicine and Biology Soc. (EMBC),* pp. 3422-3425, 2010.
[48] P.S. Tsang and V.L. Velazquez, "Diagnosticity and Multidimensional Subjective Workload Ratings," *Ergonomics,* vol. 39, no. 3, pp. 358-381, 1996.
[49] S. Rubio et al., "Evaluation of Subjective Mental Workload: A Comparison of SWAT, NASA-TLX, and Workload Profile Methods," *Applied Psychology,* vol. 53, no. 1, pp. 61-86, Jan. 2004.

Dane Powell received the BS and MS degrees in mechanical engineering from Rice University in 2008 and 2010, respectively, and has also studied at the University of British Columbia and Queen Mary, University of London. He is currently working toward the PhD degree at Rice University in the Mechatronics and Haptic Interfaces Lab. In 2010, he was awarded the NASA Space Technology Research Fellowship. He is a student member of the IEEE.

Marcia K. O'Malley received the BS degree in mechanical engineering from Purdue University in 1996, and the MS and PhD degrees in mechanical engineering from Vanderbilt University in 1999 and 2001, respectively. In 2001, she joined the Mechanical Engineering and Materials Science Department at Rice University, where she is currently an associate professor and directs the Mechatronics and Haptic Interfaces Lab. She holds a joint appointment in computer science at Rice, and is an adjunct associate professor in the Departments of Physical Medicine and Rehabilitation at both Baylor College of Medicine and the University of Texas Medical School at Houston. Additionally, she is a cofounder of Houston Medical Robotics, Inc. At Rice, her research addresses issues that arise when humans physically interact with robotic systems, with a focus on training and rehabilitation in virtual environments. In 2008, she received the George R. Brown Award for Superior Teaching at Rice University. Prof. O'Malley is a 2004 US Office of Naval Research Young Investigator and the recipient of the US National Science Foundation (NSF) CAREER Award in 2005. She received the Best Paper Award at the 2011 IEEE World Haptics Conference in Istanbul, Turkey. She is the former chair of the IEEE Technical Committee on Haptics and was on the founding editorial board of the *IEEE Transactions on Haptics*. She currently serves on the editorial board of the *ASME/IEEE Transactions on Mechatronics*. She is a member of the IEEE.