THE VIRTUAL TEACHER

R. Brent Gillespie[1], **M. Sile O'Modhrain**[2],
Philip Tang[1], **David Zaretzky**[1], **Cuong Pham**[1]

[1]Department of Mechanical Engineering
Northwestern University
Evanston, Illinois 60208
b-gillespie@nwu.edu

[2]Center for Computer Research in Music and Acoustics
Stanford University
Stanford, California 94305
sile@ccrma.stanford.edu

ABSTRACT

This paper introduces the virtual teacher, a device or agent that supplements an environment in order to facilitate acquisition by a human user of a manual skill. Like the virtual fixture, a virtual teacher generally acts as an aide or facilitator to task execution, but unlike the virtual fixture, the virtual teacher is present only during training periods. During eventual task performance the teacher is absent. The virtual teacher's objective, implicitly understood by the user, is to promote independent mastery over the task. We review and organize common paradigms for the teaching of manual skills in real-world settings and use these as inspiration for the design of virtual teachers. In particular, we are interested in the ways in which a teacher, real or virtual, can demonstrate a strategy or impart a 'feel' for a task by guiding movement of the pupil's hand. A pilot study involving 24 participants was used to test the virtual teacher concept with a simulated crane moving task. The present virtual teacher implementation did not significantly improve learning curves. However, further performance interpretations indicate that the lack of positive effect can be remedied with modifications to the virtual teacher that address component skills and ensure suitability to various initial skill levels.

1. INTRODUCTION

When instructing a pupil on the performance of a manual task, a teacher will occasionally resort to demonstration through a direct mechanical contact. Unlike visual demonstration or verbal instruction, mechanical or haptic demonstration is designed to communicate directly to the pupil's hand. Common examples of this kind of teaching occur in sports and music instruction.

With the development of haptic interfaces to virtual and telepresence environments comes an opportunity to go beyond merely providing environments in which to train users. It should also be possible to demonstrate manual strategies and perhaps even impart a "feel" for a manual task during special training periods. Note that the development of virtual environments has primarily been motivated by their postulated utility for operator training [1]. Virtual environments usually provide cheaper and safer settings for practice than the real-world settings after which they are modeled. Indeed, support now exists for claims regarding enhanced learning with haptic feedback. For example, project GROPE at UNC [2] indicates that a haptic interface can enhance learning of the docking properties of molecules. The virtual teacher concept goes beyond simply providing a training environment, it aims to provide a training environment outfitted with skilled interactive agents, ready to suggest manipulation strategies and to provide enhanced performance feedback. Our working hypothesis is the following: by using a virtual teacher to guide a pupil's hand through a task, we believe that learning times can be reduced and method recall can be improved.

The virtual teacher can be considered an extension of the virtual fixture. A virtual fixture is a mechanical impedance, appropriately modulated by the motorized interface in space or time to inspire in the operator a haptic image of a particular object. Rosenberg [3] classifies virtual fixtures as "perceptual overlays": images of objects placed on top of the user's already formed image of the primary objects. Rosenberg [3] and Sayers and Paul [4] have demonstrated that virtual fixtures can be used to enhance operator performance. The kinds of task-specific assistance provided by a virtual fixture include guiding motion with fences and preventing obstacle collisions with protective barriers. An operator may find it useful to incorporate a virtual fixture into novel strategies and thus improve task completion times.

Our virtual teacher is a perceptual overlay whose real-world counterpart is an animate teacher or coach rather than an inanimate object. We are interested in using a virtual teacher in the acquisition of sensorimotor skills and exclude intellectual or problem solving skills. We must also acknowledge that much manual skill acquisition results through visual observation of teacher demonstrations. We are only interested in learning that takes place through haptic communication between teacher and pupil. Note that teaching through demonstration and observation is not excluded when haptic communication is used alone. Like a tennis coach who grasps a pupil's wrist to impart a 'feel' for a new tennis stroke, the virtual teacher may take hold and demonstrate while the pupil observes.

The primary feature that distinguishes a virtual teacher from a virtual fixture is that the virtual teacher is present only during training periods. The goal of the teacher is to become obsolete as soon as possible, leaving the pupil to perform the skill on his or her own. Even a passive virtual object can function as a virtual teacher if the operator expects it to disappear at the end of the training period.

Virtual teachers that are active provide an even richer set of teaching paradigms. Motion paths can be demonstrated while the pupil monitors passively. Haptic encounters with the teacher can be made surprising or disconcerting if teaching by "negative feedback" is deemed useful in a particular situation. Moreover, the teacher's assistance can be withheld for periods during training to allow the pupil to "try out" the teacher's strategy. The reintroduction of the teacher then acts either to correct the pupil if their attempts have been unsuccessful, or to reinforce their learning of a successful strategy. There also exists the possibility to control the amount of assistance provided by the teacher during the training period. Krebs and Hogan's [5] Robot-aided Neuro-rehabilitation project at MIT is exploring the idea of progressively reducing the amount of "coercion" provided by a planar robot to help stroke patients regain motor skills.

Though most manual tasks can be successfully completed in a variety of ways, there is usually only one optimal way to perform the task when considering criteria such as the minimization of energy or time. McRurer [6] and Yamashita et. al.[7] have shown that a human operator, with practice, will eventually discover and adopt the optimal control strategy. This suggests that if an operator is shown the analytically obtained optimal control early on they can bypass some of the usual practice time.

A study by Repperger [8] tested the utility of active feedback in learning of a tracking task. This study showed no significant learning difference between the feedback versus non-feedback modes. In our experiment (see section 4 below) we select a task with a longer learning time and test at the end of each training period to obtain a more complete picture of our pupils' learning curves.

The remainder of this paper establishes further basis for the design of virtual teachers. In the following section, we review some recent literature on the phenomenology of presence, noting that the pupil in a teacher-outfitted virtual environment experiences not only "being in" a simulated world but also "being with" a simulated teacher. In section 3, we examine teacher-pupil interaction between humans performing manual tasks and begin to conjecture on the promise of the virtual teacher concept. We attempt to classify the various approaches a teacher can take to manually communicate skill to a pupil. In section 4, a pilot study is described in which the incorporation of a virtual teacher into a crane is tested for its utility in reducing training times. Results are presented in section 5 and discussed in section 6.

2. PHENOMENOLOGY OF REAL AND VIRTUAL PERCEPTUAL EXPERIENCE

Though it seems we live in the physical world, it is actually more descriptive to say that we each live in our own perceptual worlds. Each of us has a perceptual system which builds a mental model of our surroundings and it is according to this model that we act. This mental model is an abstraction; it comprises only that information which we pick up through our sense organs. It is so functional a representation that we often overlook the fact that interaction with the physical world is indirect, that it is mediated by our senses and perceptual systems.

Though stimulation of sense organs occurs within the boundary of our body, we attribute this sensory experience to objects in the world outside of our body. This phenomenon, an innate human tendency, has been called "distal attribution" [9]. It is distal attribution that compels us to accept our band-limited sensory impressions of the world as accurate representations of physical reality. White has points out that distal attribution is most likely to occur when an individual's sensory inputs (afferents) are lawfully related to his motor or communication outputs (efferents). Presumably, it is the individual's recognition of this lawful relationship, often contingent on related past experience, that promotes the assignment of object identity and external object location.

Owing to the unremitting operation of distal attribution, the separation between the perceptual and physical worlds is hardly ever noticed —except in rare cases prompted by perceptual illusions. Virtual reality, however, offers a contrasting case, where the phenomenology of distal attribution can more easily be recognized. As noted by Loomis [10], successful interaction with a virtual or telepresence environment hinges upon whether the operator experiences distal attribution. Indeed, a sense of presence is most likely to occur when the efferent commands issued by the operator map in a meaningful way to afferent sensory feedback from the remote or virtual environment. Virtual environment designers, then, may use the theory of distal attribution as a design directive or foundation for design methodology. Not only must the sensory experience presented to the user be suggestive of a recognizable (previously encountered) object, but the sensory experience associated with the object's response to the user's manipulations must also be suggestive of interaction with this recognizable object. Then the whole machinery of distal attribution can be expected to kick in, with a concomitant user sense of immersion.

Let us take this discussion a step further, to cover the case when not only virtual objects, but also virtual agents are present in the virtual environment. Imagine for a moment a telepresence environment in which not one but two operators cooperate to perform a task. Further, let us imagine that one operator is knowledgeable about the task at hand and wishes to demonstrate successful strategies to the other, to take on the role of a teacher. If this were a real-world situation, the pupil would have no difficulty in separating the incoming stimuli into those attributable to the task-objects and those attributable to the behaviour of the teacher, even if stimuli were restricted to the haptic modality. In a telepresence environment, however, where stimuli are filtered and often distorted, the possibility exists for attributing behaviour to the wrong distal objects. Response of a task object might incorrectly be attributed to the teacher's actions. Herein lies the challenge to designing the virtual teacher. If the teacher is to be successful, the pupil must be able to distinguish its actions from those of the task-objects in the environment. There must exist some cue or set of invariants which allow the pupil to monitor the teacher's actions independently of the responses of the object being manipulated (or co-manipulated).

It is possible that a virtual teacher which guides the pupil through a motion path by moving the object while the pupil monitors the object passively might actually undermine the pupil's sense of immersion. Gibson [11] and Katz [12] have noted that subjects are much more apt to attribute haptic properties to an object when allowed to actively explore that object than when the object is moved across their passive hand. Under passive stimulation, observers tend to describe the haptic sensations as they are felt at the skin, signifying that distal attribution has broken down. Yet in situations where a teacher and pupil are mechanically coupled through a task-object, we will occasionally want to rely on passive monitoring by the pupil. So long as real-world examples exist of successful haptic recognition by pupils of object attributes separated from teacher attributes, there exists a basis for the design of a virtual teacher that may be differentiated from virtual objects. We need only identify the relevant mechanisms at work in the real world to ensure a successful virtual world design.

The need thus arises to provide additional cues to inspire pupil awareness of the teacher's involvement. For example, if the teacher's influence contrasts strongly with that of the object alone and is present only occasionally, separation of teacher/object cues will be simpler. Alternatively, a pupil may recognize separation if the dynamics of the object and teacher are significantly different. It is possible to attribute non-passive (energy introducing) behavior to the teacher if the object being manipulated is known to be inanimate and passive. Thus, so long as the context can be established, a kind of distal attribution of the teacher can occur, where the qualities attributed are those of another human being or a fixture whose purpose is to provide training cues. Just as teachers use many modes of communication simultaneously in their work, establishing context for the virtual teacher might require special instructions to the pupil or supplemental cues in other sensory modalities.

Further, we note that teachers, in the course of imparting a manual skill, are often interested in the promotion of alternate sites of proprioception. We observe that when a teacher grasps a pupil's wrist and leads it through a motion path, the pupil occasionally becomes aware of his hand as an implement, and his forearm as the end of his arm. The teacher's actions become the pupil's sensory signals at the wrist and the pupil feels his hand being thrown around. Once the teacher has let go, the pupil attempts to throw his hand around with his forearm such that his hand feels as it did under the teacher's control. As is often the case in piano instruction, the teacher uses this method to promote a relaxed hand or a more secure coupling between forearm and fingertips.

3. TEACHER/PUPIL INTERACTION PARADIGMS

Certain conceptions of motor learning [13] hold that there are three major mechanisms by which learning takes place: method selection, chunking, and component strengthening. A teacher can get involved only in the first two methods: the teacher can point out superior methods for task completion, and the teacher can suggest superior ways to chunk the task into smaller sub-tasks. But the last method, component strengthening, relies on practice on the part of the pupil. A teacher can, however, suggest effective practicing strategies, helping the pupil avoid time consuming setbacks. In preparation for creating effective virtual teachers, we analyze certain real-world modes of teaching.

3.1 Teachers in the Real World

There exist three basic arrangements of mechanical contact between a pupil's hand, a teacher's hand, and a task object or implement handle. We use these three contact paradigms to distinguish modes by which a teacher may manually communicate a procedure to a pupil. Figure 1 shows these three modes and the paragraphs below elaborate.

Indirect Contact Paradigm (I). In the first paradigm, the teacher and pupil grasp separate points on the implement handle. There is no direct contact between teacher and pupil. The teacher wields the tool handle, and hopes that the pupil will later be able to reproduce the demonstrated

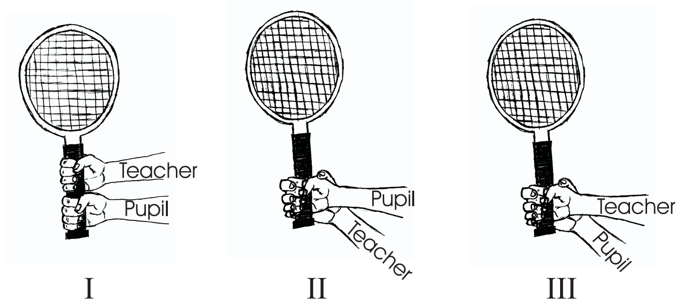

Figure 1: *Three teaching paradigms, distinguished by the arrangement of mechanical contact between teacher's hand, pupil's hand, and implement handle. I) Indirect Contact Paradigm. II) Double Contact Paradigm. III) Single Contact Paradigm.*

motion. The pupil feels a composition of the dynamics of the object and teacher. The pupil will likely find it difficult to separate the behavior of the implement from that of the teacher wielding it.

Double Contact Paradigm (II). In the second paradigm, the teacher grasps the pupil's hand which in turn grasps the handle of the implement. The pupil experiences two distinct contacts - one with the implement handle and the other with the teacher. The pupil acts as a force and motion sensor, monitoring the teacher's actions and the implement's reactions. This paradigm is the most likely candidate to promote separate distal attribution of teacher and implement on the part of the pupil, because the points of contact are distinct.

Single Contact Paradigm (III). In the third paradigm, the pupil holds the teacher's hand while the teacher manipulates the implement handle. The pupil has only one point of contact with the system. The pupil monitors the actions of the teacher and the teacher-filtered dynamics of the object. If the mechanical impedance of the object is modest compared to the impedance of the teacher's hand, the pupil will probably not be able to monitor the impedance of the object. A golf club could probably be felt through the teacher's hand whereas a piano key could not. While this is the most natural mode of manipulation for the teacher, it is probably the least likely to succeed in promoting the building of an internal model for the pupil.

3.2 Teachers in the Virtual World

In exploring the concept of the virtual teacher, we will implement the virtual equivalents to these real-world teaching paradigms. As a specific example, let us consider the task of swinging up a pendulum on a cart. One moves the cart back and forth while feeling the inertial forces of the pendulum. The objective is to move the cart such that energy is pumped into the pendulum. This task requires a fair amount of manual skill, especially when performed without visual feedback.

The virtual teacher concept suggests that, if we can include an agent that knows how hard to push and when, the pupil can learn to mimic the teacher's actions. Let us now consider how to link the teacher and pupil to the pendulum and cart in each of the three teaching paradigms differentiated by the hand/handle/hand contact arrangement. In paradigm I, the pupil is coupled to the virtual teacher through the virtual object. Note that an example virtual teacher which knows how to pump energy into the pendulum is simply a negative damper on the cart/pendulum joint. To the pupil, this system will feel unstable. A more effective teacher might be one that moves the virtual cart (interface device) back and forth at the appropriate natural frequency and in the appropriate phase relationship. To test our model for the double-contact paradigm (II) we would require two haptic interface devices, one to act as the virtual teacher on the outside of the pupil's hand, and the other to display the virtual pendulum and cart. In this way the pupil would feel forces on the back of their hand from the teacher and the response of the cart within their grasp. In scenario III, having the pupil place their hand on the teacher's hand, simply requires us to have the haptic display drive a model of a virtual pendulum/cart while the pupil holds the display to track the teacher's motion. Due to constancy of what the pupil grasps (the haptic interface), scenario (III) will not feel much different than scenario (I). The pupil will not be able to feel the difference between holding the object or holding the teacher.

4. METHOD

The pilot study presented here was designed to test the hypothesis that a virtual agent present within the workspace of a real environment could "teach" a human operator the optimal way to perform the crane-moving task. This task was chosen for its relative difficulty. We sought a task which took on the order of 5-10 minutes to provide us an opportunity to track learning curves.

Participants. Each participant was randomly assigned to one of three groups. Group A, the control group, had no assistance from the virtual teacher at any time. Of its 7 male and 1 female members, 6 were right-handed. Group B experienced the single-contact virtual teacher (scenario III). Of its 5 male and 3 female members, 7 were right-handed. Group C experienced the double-contact virtual teacher (scenario II). Of its 5 male and 3 female members, 6 were right-handed. Participant ages ranged from 18 to 40.

Apparatus. The experimental apparatus, shown in

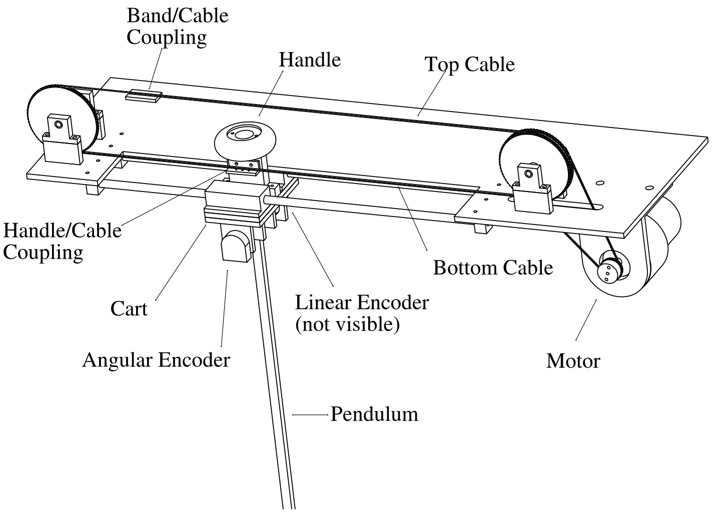

Figure 2: *Diagram of apparatus.*

Table 1. Apparatus Parameter Values

Cart Mass	0.87 kg
Pendulum Length	0.9 m
Pendulum Mass	0.011 kg
Pendulum Natural Frequency	3.58 rad/sec
Pendulum Damping Ratio	0.004

Figure 2, consists of a free-swinging pendulum pivoted to a cart and a computer controlled motor usable for driving the motion of the cart. The cart moves on linear guide rails with a 19 inch range of motion. Optical encoders sense the linear displacement of the cart and the angular displacement of the pendulum arm. Additional pertinent parameter values are indicated in Table 1. The pendulum period and damping ratio were identified by fitting a linear second order impulse response to a 100-second recording of the pendulum's response to a strike.

The cart may be coupled to the motor in one of two ways. The bottom cable may be attached to the cart handle through the handle/cable coupling (see Figure 2). Alternatively, a band worn around the participant's dominant hand may be coupled to the top cable through velcro while the participant grasps the handle. These two means of driving the cart through the motor are the single contact (III) and double contact (I) teaching modes, respectively. Group A used the same configuration as group B so that all subjects had to move against the inertia of the motor.

Procedure. Before beginning the experiment, each participant was instructed to stand facing the apparatus and grasp the handle with his/her dominant hand, palm coincident with the top face of the handle. The participant was then instructed to move the load (end of pendulum) 10 inches, starting from a state of rest (no pendulum swing) and ending in a state of rest, as quickly as possible for 50 trials. The criterion for "state of rest" allowed a small amount of oscillation (0.75 inches) in displacement of the load with respect to a target line drawn on the floor. For groups B and C, trials 11-20 and 31-40 were used as demonstrations from the virtual teacher. All other trials were unassisted. Group A performed 50 consecutive trials unassisted. Each participant wore the velcro band around his/her dominant hand during all trials.

Design of the Virtual Teacher. The virtual teacher was designed around a crane control strategy based on the *command input preshaping* technique developed and espoused by Singer and Seering [14]. The teacher essentially "knows" about the second order dynamics of the pendulum. After injecting energy into the pendulum oscillations with an initial move of the cart, the controller carefully times and sizes a second move so as to remove the previously injected energy.

Command input preshaping may be explained as follows. Figure 3 shows the response of a second order linear system to an impulse applied at t=0. Also shown is the response to an impulse of a particular magnitude applied at time T/2, where T is the damped period of the second order dynamics of interest. The second impulse magnitude is a function of the damping ratio and is designed to produce a response which cancels the first. The response to both impulses (shown with circles in Figure 3) by superposition has zero magnitude after time T/2.

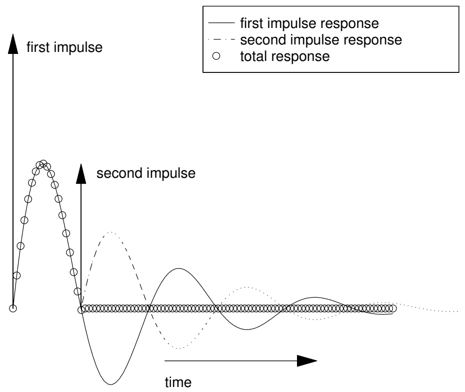

Figure 3: *Two impulse responses add to form an output which has no oscillation after the time of application of the second impulse.*

The impulse sequence thus derived is used to generate a smooth input with the same vibration canceling properties through convolution. We convovled this two member impulse train with the rising portion of a sinusoid to produce the virtual teacher's position command to the cart. The position of the cart is in turn under PD feedback control. Note

Table 2. Performance by Group and Trial Set

Group	Trials	Mean	Std Dev
A	1-10	17.3	16.3
	11-20	10.3	11.5
	21-30	7.3	4.4
	31-40	7.7	7.0
	41-50	6.6	3.9
B	1-10	20.2	15.3
	21-30	16.0	14.5
	41-50	9.4	7.6
C	1-10	18.5	17.4
	21-30	15.8	14.6
	41-50	12.8	10.8

that the pendulum oscillations are controlled in a strictly open-loop fashion. Although the cart and pendulum system is nonlinear, the oscillation amplitude of the pendulum was generally less than 15 degrees and this linear system-based technique worked quite well. Extensions to the command input preshaping technique which use longer impulse trains for robustness were not implemented in the present virtual teacher.

5. RESULTS

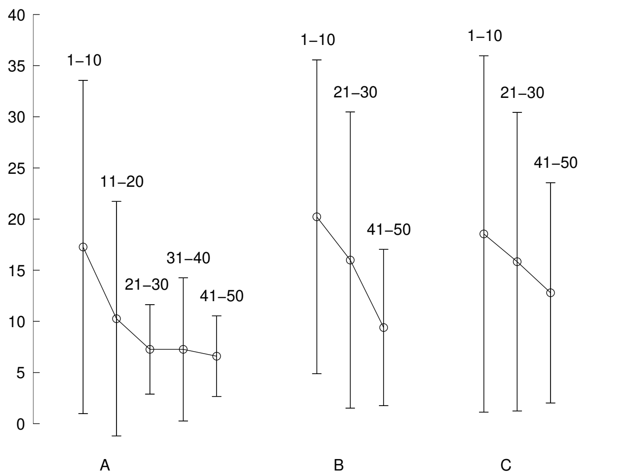

Figure 4: *Error Bar plot showing means and standard deviations over participants in each group but separated by trial sets. A: control, B: single contact teacher, C: double contact teacher.*

All performance times were separated into sets of 10 consecutive trials. Figure 4 and Table 2 show the mean times and standard deviations for each trial set in each participant group (all teacher assisted trials were omitted).

By observation of performance during the experiment and also from a look at the recorded cart and pendulum position data, we were able to group the strategies employed by participants into two basic categories. We called these the 'smooth' and the '2-step' strategies. The 'smooth' strategy is characterized by an attempt to prevent oscillations from arising by moving the cart with minimum acceleration. The '2-step' strategy is the optimal strategy employed by the virtual teacher. Figure 5 shows an example cart and load trajectory for a particular participant's implementations of the 'smooth' and '2-step' strategies. Although all participants who had been introduced to the '2-step' strategy by the virtual teacher at least attempted it, some subjects (6/16) returned to the 'smooth' strategy.

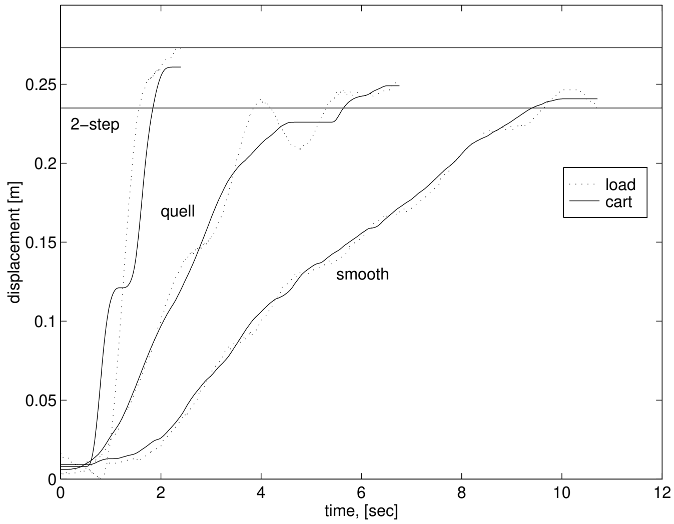

Figure 5: *Cart and load positions for three sample trials demonstrating the smooth strategy, quelling skill, and 2-step strategy. Data of a group B participant, trials 5, 30, and 44, respectively. See also Figure 6.*

We observed that a third characteristic was either present or lacking in the participant's actions during a trial. This characteristic we deemed an actual component skill and called it 'quelling'. 'Quelling' refers to motions of the cart that were produced to dampen oscillations of the pendulum. Basically, motions of the cart in the direction of the pendulum excursion would take out energy from the pendulum. For example, the center pair of traces in Figure 5 shows a cart motion which initially produced significant load oscillations yet was later accompanied by corrective cart movements which dampened these oscillations. 'Quelling' skill was requisite in the successful completion of any strategy if oscillations remained.

The first two rows in Table 3 show the number of participants in each group who attempted the 'smooth' or the '2-step' strategies. The last row in Table 3 shows the number of participants in each group who, to our best judgement,

Table 3. Strategies Attempted and Quell Skill Mastery

	Participant Group		
	A	B	C
Attempted 'Smooth'	8	5	7
Attempted '2-Step'	4	8	8
Mastered 'Quell'	4	2	2

mastered 'quelling'. It was actually not difficult to determine which participants possessed or acquired the quell skill and which did not.

Figure 6 shows the learning curve for the same example participant from group B whose data are shown in Figure 5. The strategy employed in trials 1-10 was 'smooth'. Once the virtual teacher demonstrated the '2-step' for the participant, however, this new strategy was attempted for about half of the next 10 trials and performance clearly worsened. But after the second session with the teacher, only the '2-step' strategy was used and performance was roughly similar to the first 10 trials. But note that the best performance attained using the '2-step' strategy was superior to that attained using the 'smooth' strategy. That best unassisted performance compares to the assisted performance seen in Figure 6 for trials 11-20 and 31-40.

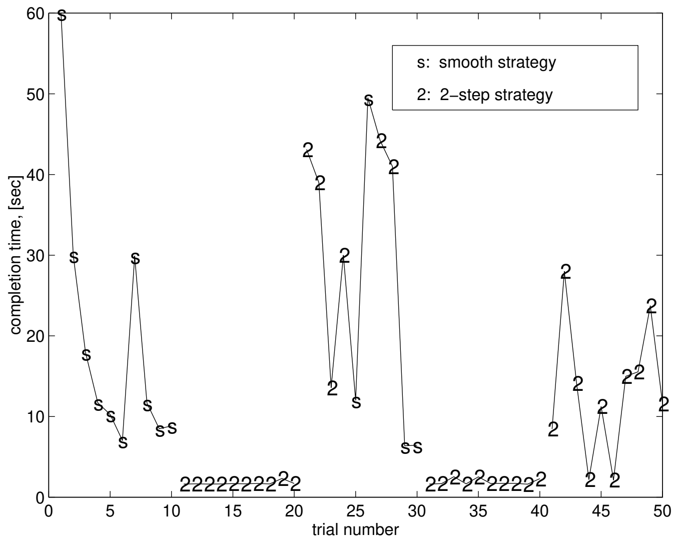

Figure 6: *Learning curve for the same group B participant of Figure 5.*

6. DISCUSSION

Clearly from Figure 4, proficiency is being acquired in all three groups. However, from the averaged learning curves, there is no conclusive evidence to indicate a positive effect from the virtual teacher. In fact, comparing group A to the teacher groups B and C suggests that the teacher had a negative effect on performance. However, there exist a number of observations to explain this phenomenon and indicate that the teacher may easily be remedied to yield better results.

Comparison of control group A against groups B and C. As quickly became evident during the experiment, and as the data also show (see the last row of Table 3,) most participants could have profited from a teacher who instructed on the 'quelling' technique. In case any oscillations remained after an attempted '2-step' move (which would occur given inaccuracies in timing or sizing,) proficiency in 'quelling' was required. From the perspective of many participants, the '2-step' teacher was too advanced. One of these participants commented, "It's like watching Michael Jordan". The virtual teacher made the task look easy. Without the component skills, learning from an advanced teacher can be difficult. Thus our virtual teacher needs to broaden its curriculum. We believe that a teacher who simply moved the cart in the direction of the pendulum excursion would have been quite effective for teaching the 'quell' skill. A feedback controller could easily be developed for this purpose.

We did observe that all participants who learned how to quell oscillations (4/16) were able to successfully adopt the '2-step' strategy. Some participants (4/16) only partially solved the problem by mimicking the virtual teacher's '2-step' strategy. Most notably, of all participants who mastered the quell skill and attempted the '2-step' strategy, none returned to the smooth strategy.

At this time, our results are inconclusive as support for virtual teacher. As seen in Figure 4, the downward trend over trial sets in the mean performance evident in groups B and C was not significantly steeper than for group A. In fact, groups B and C did not show the diminishing size in standard deviation as trials sets progressed which was present for group A.

These trends can be explained by the change in strategy that invariably followed the virtual teacher demonstrations. Because the '2-step' was quite an aggressive strategy and rather sensitive to timing and move-sizing variations, participants often found themselves in high amplitude oscillation situations following an attempted '2-step'. We presume that this ultimately led to performance degradation because the 'quelling' skill was not available for many participants to use in combination with the '2-step' strategy. The data from the example participant in group B confirm this hypothesis (See Figure 6). We conjecture that our results would have given stronger support for the virtual teacher had a teacher been available for instruction in the 'quelling' skill. Future iterations on the experiment will include teaching algorithms designed to teach 'quell' and perhaps will allow participants to choose when they want the teacher's help, or to choose from among several available teachers.

Comparison of groups B and C. The particular apparatus used for our present experiment featured very little haptic feedback. The mass of the cart and coupled inertia of the motor and bearing friction dominated the inertia forces of the pendulum. Thus our participants could not monitor the motions of the pendulum by feel, rather they were faced with what was basically a hand-eye coordination task. We theorize that this dependence on visual rather than haptic feedback explains the lack of differentiation between groups B and C. Changing the focus of the task to be more haptically-oriented may yield better data with which to contrast the effects of the single contact and double contact teachers. Adding a large weight to the end of the pendulum and perhaps using a curtain to mask the whole pendulum may be useful.

Insights and Anecdotal Findings. One solid result of our pilot study was that the virtual teacher was indeed able to effectively demonstrate and encourage the adoption of an alternate strategy. Although only 4 participants in group A attempted the '2-step', all participants in groups B and C at least attempted the optimal '2-step' strategy (See Table 3.) The virtual teacher was able to non-verbally communicate the essentials of the optimal strategy. Interestingly, all but three subjects in groups B and C could describe the '2-step' strategy in their own words, even if they had not yet mastered it.

The responses to our post-experiment questions were quite informative. In groups B and C, all most all participants supported the idea of verbal instruction complementing (perhaps even replacing) the virtual teacher. Surprisingly, though, only a slight majority replied that a human teacher would have been more effective (9/16), mostly attributing this to the human's ability to verbally instruct. Among those who disagreed, almost all commented that for demonstrating motor skills, the virtual teacher has an advantage over the human teacher because of its accuracy and consistency.

7. CONCLUSION

We are interested in a virtual environment that, in addition to being populated with virtual objects, contains a virtual teacher that can demonstrate certain object manipulation techniques. We have drawn on the concept of distal attribution to theorize about the communication of manual skill from teacher to pupil. A pilot study with 24 participants used the crane moving (cart and pendulum) task to test the ability of a virtual teacher to communicate the time-optimal strategy. Although average performance time for the experimental groups did not improve faster than that of the control group, the optimal strategy was successfully communicated to the experimental groups. Our pilot study points to the need for an accompanying virtual teacher to teach the component skill of quelling load (pendulum) oscillations.

REFERENCES

[1] N. Durlach, *Virtual Reality: Scientific and Technological Challenges*, vol. National Research Council. National Academy of Sciences, 1995.

[2] F. P. Brooks Jr., M. Ouh-Young, J. J. Batter, and P. J. Kilpatrick, "Project GROPE-haptic displays for scientific visualization," in *Computer Graphics*, vol. 24 of *4*, pp. 177–185, ACM, August 1990.

[3] L. B. Rosenberg, "The use of virtual fixtures to enhance teleoperator performance in remote environments with time delay," in *Haptic Interfaces for Virtual Environments and Teleoperator Systems, ASME WAM*, (New Orleans), 1993.

[4] C. Sayers and R. Paul, "Synthetic fixturing," in *Advances in Robotics, Mechatronics and Haptic Interface (DSC-Vol.49)*, (New Orleans, LA.), pp. 37–46, 1993.

[5] H. I. Krebs and N. Hogan, "Robot aided neuro rehabilitation (web page)," in *http://me.mit.edu/groups/hogan /res/neuroperf.html*, 1998.

[6] D. McRuer, "Human dynamics in man-machine systems," *Automatica*, vol. 16, pp. 237–253, 1980.

[7] T. Yamashita, T. Nakamura, M. Ohtsuka, and T. Taniguchi, "Observation and analysis of heuristic learning process in manual control of a crane," *Digital Systems for Industrial Automation*, vol. 2, no. 1, pp. 31–48, 1983.

[8] D. W. Repperger and C. Goodyear, "Active controllers and time duration to learn a task," in *Conference on Manual Control*, 1985.

[9] B. W. White, F. A. Saunders, L. Scadden, P. Bach-y Rita, and C. C. Collins, "Seeing with the skin," *Perception and Psychophysics*, vol. 7, pp. 23–27, 1970.

[10] J. M. Loomis, "Distal attribution and presence," *Presence*, vol. 1, pp. 113–119, Winter 1992.

[11] J. J. Gibson, "Observations on active touch," *Psychological Review*, vol. 69, pp. 477–491, 1962.

[12] D. Katz, *The World of Touch [Aufbau der Tastwelt (1925)]*. Lawrence Erlbaum Associates, 1989. Translated by Krueger, Lester E.

[13] S. W. Keele, *Motor Control, in: Handbook of Perception and Human Performance*. John Wiley, 1980.

[14] N. C. Singer and W. P. Seering, "Preshaping command inputs to reduce system vibration," *Dynamic Systems and Controls*, vol. 112, pp. 76–82, March 1990.