Handbook of Virtual Environments: Design, Implementation, and Applications
Edited by Kay M. Stanney, Lawrence Elbraum Associates, 2002

Chapter 49

Use of Virtual Environments in Motor Learning and Rehabilitation

Maureen K. Holden, Ph.D., P.T
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology, Cambridge, MA 02139;
Cambridge, MA 02139
USA
Phone: 617/ 253-5972; Fax: 617/ 258-5342; e-mail: holden@ai.mit.edu

and

Emanuel Todorov, Ph.D.
Gatsby Computational Neuroscience Unit, University College London
17 Queen Square, London WC1N 3AR
UK
Phone: +44 (0) 20 7679 1185; Fax: +44 (0) 20 7679 1173
email: emo@gatsby.ucl.ac.uk ; web: www.gatsby.ucl.ac.uk/~emo

1. Introduction

The purpose of this chapter is to describe how virtual environments (VE) can be utilized to facilitate motor learning in normal subjects, and to enhance the rehabilitation of disabled subjects.

The chapter begins with a review of relevant literature on the use of VE for training sensori-motor tasks in normal populations in two main areas – spatial-motor and navigation learning studies. Next, clinical literature on the potential and actual applications of VE to the field of rehabilitation is reviewed. Advantages and disadvantages for the use of VE, discovered through these studies, are identified and discussed.

We next describe a system we have designed, whose purpose is to provide a new method for motor learning, 'by imitation', within a virtual environment. A number of key design issues that were encountered during development of the system, along with the solutions chosen to address them are described. Finally, we provide summaries of results to date based on using the system to train normal subjects and patients with stroke on selected motor tasks. We end with a short description of ongoing work and future directions.

2. Literature Review: Motor Training using VEs in Normal Subjects

In recent years, there has been great interest in using VE for learning various types of motor and spatial tasks (Darken, Allard and Achille, 1998; Kozak, Hancock and Arthur, 1993; Regian, Shebilske and Monk, 1992). This interest has been driven by a number of perceived advantages of virtual over real-world training. These advantages are both practical and theoretical in nature. Practical advantages are discussed below; theoretical advantages are discussed in the context of system design in section 4.1.

2.1 Practical advantages of VE training

Practical advantages of using VE for training include safety, time, space and equipment, cost efficiency and documentation considerations: 1) Safety - In practicing a virtual task which is potentially dangerous (such as flying a plane for a normal subject; or lifting a pan of hot water for a disabled subject) there are no negative consequences to failure; 2) Time - The training task can (in theory) be quickly altered to become easier or more difficult. 3) Space and Equipment - Requirements may be minimal compared to the real environment (for example, training military personnel on a 'virtual' vs. real ship); 4) Cost Efficiency - Possibility for fewer training personnel (or therapists in the case of patients), while still maintaining a similar outcome level following training or treatment. This may be achieved by semi-automating training on the computer, allowing subjects to work independently some of the time. In medical applications, telemedicine has the potential to reduce the cost of providing care to remote areas, while improving access for remote areas to higher quality of care; 5) Documentation - Automatic scoring and report systems can be developed to allow easy monitoring of progress of subjects using the VE training system.

Obviously, the potential to apply VE for training to all sorts of sensorimotor tasks in fields as diverse as defense, medicine, industry, and sports, is huge (witness the size of this volume!). In looking at the literature on VE applications to motor learning – broadly defined – in normal subjects, one finds that much work is geared toward very complex motor skills which have a large cognitive component (such as military training for flying planes, or navigating ships (Kennedy, Lanham, and Massey, et.al.,1995; Lawson, Rupert, Guedry et al, 1997; Theasby, 1992)), or a very high level of eye-motor coordination (such as VE training aimed at virtual surgical training for physicians (Docimo, Moore and Kavoussi, 1997). The issues important to researchers studying these types of problems are often quite different from those faced by the researcher whose goal is to develop a system that would be useful in helping disabled subjects relearn simple movement skills like lifting a cup, or turning a key in a lock, or to train skills that are more predominantly perceptual-motor in nature, such as a tennis swing.

In fact we found relatively few studies in the literature on using VE's to train such tasks. The most relevant studies fell into two classes 1) studies of spatial-motor learning such as the 'pick and place' task (Kozak, Hancock, and Arthur, et al., 1993), or procedural console-operations tasks (Regian, Shebilske and Monk, 1992); and 2) studies of learning to navigate in small (Waller, Hunt and Knapp 1998) or larger scale environments (Wilson, Foreman and Tlauka, 1997; Ruddle, Payne and Jones, 1999).

2.2 Spatial-Motor Learning in VE

Kozak et al. (1993) trained subjects to move a series of cans, positioned in a line from left to right directly in front of them, to matching discs placed 6 in. forward of the start location. The sequence was then repeated in reverse order, so that the cans ended up at the original starting location. Three groups were used; one group trained in a virtual environment, one trained on the real world task, and one group received no training, but only testing. Subjects who trained in the real world were significantly better than either the VE or no training groups. And subjects who trained in VE were no better at performing the real world task following training than were subjects who received no training.

However, subjects who trained in VE *did* improve their performance *during* VE, at a rate comparable to subjects who trained in the real world. That is, they *learned* the virtual task, but learning the virtual task did not *transfer* to improved performance on the comparable real world task. Proposed reasons for the lack of transfer included subtle differences in the grasp required in the two situations (VE vs. Real World), system time lag, differences in the field of view in which the hand could be seen, lack of tactile and acoustic feedback, difficulty aligning some subjects with different body sizes so that all the virtual cans could be reached. If subjects were using the method of loci (Bower, 1972) as a learning strategy (in this method, subjects associate motor responses with imaged spatial locations), then slight differences in displayed locations relative to the subject's arm in the two conditions could have interfered with learning, and transfer of skill to the real world test.

A later study by Richard, Birebent, Coiffet et al. (1996) examined the influence of several of the factors identified by Kozak et al (1993) as possible factors that interfered with transfer of learning. Richard et al (1996) examined the effect of different graphics frame rates and viewing modes on grasp acquisition time, and the effect of force feedback and added acoustic signals on force regulation ability during grasp in a virtual environment. They found no difference in object capture times (time to grasp a virtual object) between stereo and mono displays for frame rates of 28 or 14 frames/sec (fps). However, when frame rate fell below 7 fps, stereo viewing increased performance by 50%. (Note that currently available graphics hardware allows higher frame rates at lower cost.) Direct haptic feedback was found to be superior to pseudo-feedback (reduced error rates and improved performance of force regulation task). The lowest error scores and task completion times occurred when direct haptic feedback was augmented by an acoustic signal which indicated the initial contact of the virtual hand with the virtual ball.

Regian et al. (1992) trained subjects in a procedural learning task (a complex 17-step console operation task) using a virtual environment, and found that all 31 subjects in the experiment were able to learn the task using VE for training. The purpose of the experiment was to compare the effects of two types of instructional strategies on learning, therefore, transfer to real world performance was not tested. However, this study did serve to establish that it was possible to learn motor skills in a virtual environment.

2.3 Motor Learning and Transfer

The most essential issue in considering the usefulness of VE training for motor control in normal subjects, as well as for rehabilitation of patients, is the ability to *transfer* the performance that is learned in the *virtual* world to performance of that same task in the *real* world. It was surprising to see that many studies examined learning in a virtual environment, but did not measure performance on a similar real-world tasks following training (Gillner and Mallot, 1998; Ruddle, Payne and Jones, 1999). Perhaps this occurs because researchers are aware of the many technical factors which account for discrepancies between the virtual and real worlds, and the goal of many studies is to understand these factors better, so that the fidelity of the virtual environments can be improved. However, without testing transfer to real world performance, much effort could be consumed in improving fidelity in VE systems beyond the level that is really necessary to accomplish effective training.

2.4 Navigational Learning

Two studies using VE to teach spatial-navigation knowledge found that under certain conditions training transferred well to the real world environment (Waller et al. 1998; Wilson et al. 1997). In the first of these studies, Waller et al. (1998) examined the effect of six different training environments on subjects ability to perform blindfolded in a real world 14x18 ft. maze environment. The conditions compared were: 1) no training; 2) real world training, 3) map; 4) VE desktop, 5) VE short-immersive (1-min.), and 6) VE long immersive (2-min). Subjects (n=120) received six consecutive bouts of training on one of these conditions. Each training bout was followed by a test trial on the real world maze. Results showed that for early learning (defined as performance on Test trials 1+2), real world training was significantly better than all other methods, and although all other conditions were equal statistically, there was a trend to *worse* performance for subjects trained in VE (x=270 s for VE subjects; 163 s for real world trained subjects). However, by the 6th trial, mean times for the real and long-immersed VE conditions were statistically equal, and significantly shorter than for the other two VE conditions (desk-top and short-immersive). Although this implies superiority for immersive VE over desktop, it is not clear that the superior performance for the immersive condition would have held if subjects had trained in the other conditions for equivalent amounts of time as in the long-immersive VE condition. In examining the issue of transfer from VE training to real world performance, this study is of interest, because it shows that both desktop and immersive VE methods were capable of improving performance on a similar real world task. In addition, when controlled for time spent in training, there was no difference in training effectiveness of the desktop and immersive VE methods.

Wilson et al. (1997) examined navigational learning in a larger environment, via exploration of a real multi-story building and a to-scale computer simulation of the same building. The computer simulation was presented to the subjects using a desktop VE system. The computer keyboard was used to navigate. A control group received no training. The results showed that subjects who received the simulation training were able to transfer the 3-D spatial knowledge learned in the VE to a variety of spatial measures in the real world. For many measures, the VE trained subjects performed as well as subjects who trained in the real environment, and for most measures the VE subjects performed better than the control group. This study is of interest because it shows that learning and transfer can occur with less sophisticated desktop systems.

2.5 Cybersickness and Negative After-Effects

Another issue that stands out in many studies of VE training performed on normal subjects is the relatively common incidence of negative side effects. These side effects include cybersickness, altered visual-motor and sensori-motor coordination, and postural disequilibrium (Cobb, 1999; Nichols, 1999; Stanney, Kennedy, Drexler and Harm, 1999; Stanney and Salvendy, 1998). These side effects have the potential to negatively affect motor coordination in the real world post training.

Such problems are of even greater concern for patient populations, especially if central nervous system function is impaired due to disease or injury. Minor alterations in postural equilibrium or eye-hand coordination that could safely be handled by a subject with a normal nervous system, might, in a patient with impaired nervous system function, result in a fall due to loss of balance or inability to catch oneself with an arm that reaches to the incorrect location in space for support. Because the topic of negative aftereffects is covered in detail elsewhere in this volume, an extensive review is not provided here. However, in the sections on design (4.0, 5.1 and 6.1) we discuss how findings from research in this area have influenced design decisions in the system we are presently using. As more is learned about the exact causes of these side effects through studies on normal subjects, design changes in future systems and improved training routines will be better able to prevent the occurrence of unwanted side effects in normal, and hopefully, disabled subjects as well.

3. Clinical Literature: Virtual Environment use in Rehabilitation

In the field of rehabilitation, VE application is still in its infancy. Much of the clinical literature published to date on virtual environments in rehabilitation has been focused on describing the potential applications of the method to different areas of rehabilitation (Greenleaf and Tovar 1993; Kuhlen and Dohle, 1995; Rose, Attree and Johnson, 1996; Wilson, Foreman and Stanton, 1997). However, relatively few reports of actual use and effectiveness of VE with patient populations have been published. While many investigators are working in this field, much of the work is still under development or just beginning to move from the development to the clinical testing phase. In the first section below we briefly describe the different areas of application that are under development to give the reader a broad sense of the field. In the second section we review several reports relevant to motor learning of upper extremity tasks in disabled subjects.

3.1 Overview of VE applications in rehabilitation

Proposed applications of VE to the field of rehabilitation fall into four major categories: 1) use in measurement/ diagnostic tests; 2) assistive technology; 3) social and entertainment applications; and 4) training of impaired functions. Training of impaired functions may be divided into training of tasks which are primarily a) perceptual-motor, or b) cognitive in nature.

3.1.1.Applications for Measurement and Diagnostic Tests. Kuhlen and Dohle (1995) describe a system designed to aid physicians in the analysis and diagnosis of movement disorders, such as different types of pareses and apraxias. The movements of patients performing different motor tasks are recorded, then the trajectories of these movements can be displayed in 3-D, viewed from different angles, and precisely quantified. VE can also be used to explore interactively medical imaging data such as MRI and PET scans of the brain, allowing physicians to gain a clearer understanding of the location and size of lesions. When paired with clinical evaluations from patients, such information should enhance the ability to predict structure-function relationships in the brain. Rose et al. (1996) propose using VE to make the neuropsychological assessments of cognitive function in patients with brain injury more ecologically valid. This could be done in VE by measuring cognitive function in the context of realistic everyday functions, such as cooking in a virtual kitchen. Another advantage of VE testing in brain injured subjects would be the ability to distinguish the contribution of cognitive vs. motor impairments in task performance. Because patients with brain injury often have both cognitive and motor impairments, these two factors are mixed together in real world testing. VE testing could allow users to focus on only the cognitive factors by reducing the motor parts of the task in VE to one simple movement.

3.1.2 *Applications for Assistive Technology.* Assistive technology in general is designed to help disabled persons use different types of technology to substitute for functional abilities they have lost. VE would be one more tool to extend the possibilities for such individuals. For example, telerobotics could be applied to allow users with very little motor control to manipulate objects in their environment by interfacing with a robot device. Warner, Anderson and Johanson (1994) are working on a system to allow quadriplegics such control. In another example, Greenleaf and Tovar (1994) describe a 'Virtual Receptionist' system that allows people with speech or motor impairments to perform the job of receptionist: answering and making phone calls, and other secretarial tasks. Two groups have worked on systems that use a VE data glove to record sign language gestures, then translate these gestures into written or spoken words, by using a speech synthesizer (Kramer and Leifer, 1989; Vamplew and Adams, 1992). Such a system would allow deaf individuals to converse freely with hearing persons who are not knowledgeable in sign language.

3.1.3 Social and Entertainment Applications. Individuals with severe disabilities often experience social isolation due to difficulties with transportation, energy levels, and access. VE has been proposed as a way of allowing handicapped individuals at remote sites to join together and interact in a virtual world (Smythe, 1993).

3.1.4 Training of Impaired Functions. In the field of locomotor training, a prototype device which can superimpose objects into the field of view of patients with Parkinson's disease, while still allowing them to see the real world, has shown promise as a method to counteract the slow shuffling gait which is characteristic of patients with this disorder (Emmet, 1994; Reiss and Weghorst, 1995). A method to teach children with orthopedic impairments to operate motorized wheelchairs has been developed by Inmman, Peaks, Loge et al. (1994). Another VE system designed to train standing posture and dynamic balance control in the elderly, is currently being tested (Cunningham, 1999). The system is based on video capture in a non-immersive environment.

A method for treating strabismus (a deviation of the eye position in the socket, often due to weakness in one or more eye muscles) by exercising eye muscles in VE has been proposed (Lusted and Knapp, 1992). By attempting to keep objects in the virtual environment aligned, the eye muscles could be exercised; the difficulty could be calibrated based on the degree of misalignment of the resting eye axis of the patient.

Finally, methods are being developed to assess and train patients with acquired brain injuries on cognitive function tasks (Christiansen, Abreu, Ottenbacher et al. 1998; Pugnetti, Mendozzi, Motta, et al., 1995). The method being developed by Christiansen et al. uses a virtual kitchen to assess performance on daily tasks, such as preparing a can of soup. The focus of the task is cognitive, as the subject uses a mouse click or joystick to navigate in the environment and accomplish each step of the task in sequence. That is, subjects are not required to move a virtual (or real) arm to actually do the task as they would in the real world. Such a test can help therapists determine the degree to which cognitive vs. motor deficits contribute to a patient's impaired function on a variety of activities of daily living.

3.2 Reports on VE training results in rehabilitation

3.2.1 Spatial Learning. Children with severe motor disabilities often have poorly developed spatial awareness (Foreman, Orencas, and Nicholas et al. 1989). One significant contributing factor to this deficit is thought to be the reduced ability to explore the environment in an active way secondary to the mobility impairment. Virtual environment technology has been proposed as a way to offer such children a chance for active, independent spatial exploration.

Wilson, Foreman and Tlauka (1996) decided to test whether such training in VE would result in improved spatial knowledge in a comparable real world setting. The authors used desktop VE system to train large scale spatial knowledge of a 2 story building. Subjects (n=10) were children with a variety of disabling conditions, including spina bifida, cerebral palsy, and muscular dystrophy. All relied on a wheelchair as their prime method of mobility, though 3 subjects could walk short distances with a walker. They ranged in age from 7 to 11 yr. A control group of 8 healthy adults, mean age of 23 yr., received no VE training, but underwent the same testing as the children. Children received one training session. They explored a virtual model of the actual building they were in for the experiment, with the goal of finding the fire extinguishers and the fire exit door. Immediately following VE training, they were tested on knowledge of the fire extinguishers and fire door locations in the real building. This was done in two ways. First, while still in the training room with the computer, subjects used a pointer to aim at the expected location. This was followed by a wheelchair tour of the real building during which subjects were asked to guide the experimenter to two specific locations. Results showed that all the subjects were able to point to and find the desired locations in the real world, and to do so significantly better than control subjects who received no training.

The results of this experiment are significant for our work, because the system used was a relatively simple VE desktop display. Subjects had no vestibular or kinesthetic feedback during their exploration of the virtual world; the feedback was purely visual in nature. Despite these limiting factors, the exposure of only one session resulted in significant learning.

3.2.2 Motor Learning. Researchers at Rutgers (Burdea, Deshpande, Liu et al., 1997) have developed an interesting line of work using a VE system designed for hand rehabilitation and diagnosis. The system uses a PC workstation, Rutgers Master II force feedback glove and a Multipurpose Haptic Control Interface, which allow measurement of grasping forces applied to 16 regions of the hand. The system was initially designed as a diagnostic tool for hand rehabilitation, then developed further to include rehabilitation capabilities. WorldToolKit was used to create exercise routines modeled after standard hand rehabilitation exercises. These include squeezing a ball, compressing ‘silly putty’, isolated finger flexion/extension, and functional tasks of pegboard insertion and ball throwing. The interaction with virtual objects is accomplished with different hand gestures, e.g., grasp, lateral pinch, pointing, release. A desktop display is used. The forces generated by the patient in the virtual movement are fed back through the glove with a mechanical feedback bandwidth of 10-20 Hz. The system is currently undergoing clinical trials at Stanford University. Telerehabilitation capability has recently been developed for the system (Popescu, Burdea, Bouzit and Hentz, 2000), and plans are underway to extend the rehabilitation capability of the system to the elbow and knee.

Several additional studies on motor learning in normal subjects (Todorov, Shadmehr, & Bizzi, 1997) and subjects with stroke (Holden, Todorov, Callahan & Bizzi, 1999; Piron, Dam, Trivello et al., 1999; Piron, Trivello, Cenni, et al., 2000), utilize the VE system we have developed. The results of these studies are therefore presented following a description of the 'learning by imitation' system, in sections 5.2, 6.2 and 6.3 of the chapter.

4. Design of a System for 'Learning by Imitation" in VE

4.1 Theoretical considerations

Observations of human movement suggest that the nervous system uses some sort of abstract motor plan, and fills in the details needed to instantiate that plan in specific situations (Schmidt, 1975). These observations fall in two categories:

- Certain geometric features of movement trajectories (the individual writing style, the shape of the speed profiles) are preserved when the movement is executed with different actuators (left or right arm, foot), requiring very different muscle activation patterns and dynamic interactions with the environment (Atkeson & Hollerbach, 1985; Laquaniti, Terzuolo, & Viviani, 1983; Morasso, 1981).
- If the movement trajectory is repeatedly distorted by external forces, or the subject perceives the trajectory as being distorted (due to distorted visual feedback), the motor system often adapts so that the (perceived) movement trajectory returns to normal (Lackner & DiZio, 1994; Shadmehr & Mussa-Ivaldi, 1994; Wolpert, Ghahramani & Jordan, 1995).

Such observations have led a number of investigators to believe that the geometric (kinematic) aspects of the movement form a major component of the abstract motor plan. If this is the case, the acquisition and rehabilitation of motor skills should be enhanced by training methods that somehow create an internal representation of the desired kinematic pattern (Schmidt & Young, 1991). An obvious way to do this is to learn by imitation: show people a detailed movement, ask them to reproduce that movement, and provide feedback on the mismatch between the desired and actual movement.

While this can be accomplished with traditional methods such as demonstration or video playback, imitation learning systems utilizing virtual environments offer a number of advantages. Some of these practical advantages were discussed earlier and are common to VE systems in general, for example, space and equipment requirements are less, and mistakes made during practice have no negative consequences. The more important advantage of using VE technology for imitation learning however, is its potential to enhance motor learning. This is due to several factors. The most important factor is the unique capability for real time feedback in the very intuitive and interpretable form that VE provides. Other advantages are that learners can see their own movement attempts in the same spatial frame of reference as that of the teacher (unlike practice with a real coach or therapist) and that the task can be simplified in the early stages of learning, allowing the learner to focus on key elements of the task. Training environments can also be customized for different therapeutic purposes and the system designed to help the learner detect and correct errors more rapidly. Finally, practice, an essential element in motor learning, can be facilitated by making the task fun.

4.2 Design consideration for normal subjects

The basic idea for implementing learning by imitation is illustrated in Figure 1. We have mounted a computer monitor at eye level, with the screen facing down. A see-through mirror is mounted 20 cm below the screen. Thus images displayed on the screen appear in a workspace below the mirror (in 3D when stereo glasses are used). The user of the system can see both his hand and the computer generated images, in the same physical space (the picture is taken from the user's point of view, while the user is looking down trough the mirror). The user can "insert" his hand in the virtual hand and attempt to match the posture exactly. The virtual hand can display an animation of a desired movement, that the user attempts to follow repeatedly. Both the animation and the real movement occur simultaneously (overlaid or next to each other), making any deviations between the two immediately obvious.

While this would be the ideal form of imitation learning, it is hard to implement such training systems because existing technology does not always permit superimposing virtual images on the real world. In cases where the workspace is large and monitors/mirrors cannot be easily mounted, we would have to use a head-mounted display to present the virtual image. That image will move with the head of the user unless head-tracking is used for correction. Our experience with head-tracking (even with highly accurate sensors such as Optotrak 3020) suggests that it is extremely difficult to fool the human visual system – the virtual objects always look unstable due to tracking/calibration errors and delays. This is perhaps not a major problem when the virtual objects are used to provide abstract information; in our case however we are dealing with a real time visuo-motor loop and a motor learning system that is used to a perfectly "calibrated" physical world. So we decided to leave the ideal implementation for the future, and focused on more feasible alternatives where both the desired and actual movements are presented in a virtual environment.

4.3 Design considerations for rehabilitation subjects

In considering the kinds of system features that would make a VE system useful for motor training in patients with neurological impairments, one can begin by asking the question, 'What is different about training patients as compared to normal subjects?'

In contrast to healthy subjects, where the focus of VE training is often complex perceptual-motor skills, such as flying a plane, in rehabilitation the goal is usually to re-learn very simple motor tasks – ones we all take for granted – and to do so in the face of a variety of sensory, motor and cognitive deficits. The exact type of deficit can vary tremendously from patient to patient, even with the same diagnosis. To accommodate for this variety, the first characteristic we need to think about is flexibility, so that the training can easily be adjusted to the needs of different patients. Another major focus is to facilitate transfer of motor learning that occurs in the virtual world to performance in the real world. Flexibility in system design helps us to ensure that the virtual tasks we create resemble the real tasks as much as possible, particularly in terms of the movement required to perform the task.

Figure 1. A photograph taken from the user's point of view, while the user is looking down through a semitransparent horizontal mirror. A monitor mounted above the mirror (facing down) projects virtual images, that appear aligned with the physical hand in the same workspace. Note that appropriate illumination has to be used in order to see both hands.

Because the deficits patients have can make learning even simple movements daunting, we need to utilize VE in a way that will allow us to *enhance* the feedback provided to the patient. The idea of augmented feedback is not a new one. But, applying it in the context of VE, in combination with learning by imitation of a virtual teacher, in real time, does differ from techniques used in the past. Thus, we have incorporated into the system a number of ways for the patient to detect the exact kinds of errors they are making during attempted movements. Simply trying to imitate a desired movement may not be enough, particularly for patients with a sensory deficit who cannot 'feel' exactly what their arm is doing. (Note, however, many of the features we have designed with patients in mind could well enhance motor learning in normal subjects as well.)

What about graphics displays? While very complex, realistic or fanciful graphics and fast-paced game formats are appealing to normal subjects for VE training, they may be overwhelming for patients. When the brain itself has been affected, as in stroke or acquired brain injury (in contrast to disabilities where the brain is intact, like spinal cord injuries), processing of all this information may be too difficult. In contrast, a very simple display, with few movements, may help a patient to focus on the task at hand, and enhance learning.

Next, we need to think about the cardinal rule in medicine: “First, do no harm.” The issue of possible negative side effects such as cybersickness, altered eye-motor coordination, and postural disequilibrium must be carefully considered. To date, there seem to be no reports of negative side effects using desktop display systems for VE. Most of the problems seem related to the immersive environment created in head mounted displays (HMD) (Cobb, 1999; Nichols, 1999; Stanney, Kennedy, Drexler and Harm, 1999; Stanney and Salvendy, 1998). It seems obvious that feelings of disorientation, postural disequilibrium, or aftereffects in arm control that produce past-pointing, could be much more dangerous for disabled subjects, in whom a fall and injury could more easily be elicited than in normal subjects.

Even without aftereffects, HMD’s present some problems. The extra weight imposed on the head will be much more difficult for patients to handle, because trunk and postural control muscles are often adversely affected by neurological impairments. Also, in the HMD, the patient cannot see the therapist or her instructions. Even with see-through HMD’s, problems with post-exposure alterations in eye-hand coordination and felt limb position have been noted (Biocca & Rolland, 1998; Rolland, Biocca, Barlow et al. 1995). If an HMD is used without a head position tracker, (and visual display is thus fixed relative to the head), the patient may become disoriented due to inaccurate cues for real world horizontal and vertical, and begin tilting to one side while sitting or standing.

Finally, practical issues are important. Keeping costs down, making systems easy to use, and minimizing the amount of equipment attached to the patient, will offer the greatest potential for success and acceptance of VE as a standard treatment procedure in rehabilitation clinics.

Because of many of the factors listed above, we have designed our system to be used with a desktop display. The display can also be projected on a large screen instead of a monitor, to give a greater feeling of immersion. A significant disadvantage of the desktop display is poor depth visualization. This can be compensated for in part by training features which allow us to playback the recorded performance of the patient, then rotate the display in 3-D in any desired direction, to show the patient where his/her errors are, and how to correct them.

4.4 System components

To date we have built and tested two virtual environments – one for training ping-pong shots (Todorov et al. 1997) and the other one for rehabilitation of stroke patients (Holden et al. 1999). We first describe the features common to both systems, followed by a more detailed discussion of each system and our experimental results. Components described in sections 4.2.1 – 4.2.3 were used in our first two experiments. Based on results in these experiments, components listed in sections 4.2.4 - 4.2.6 were developed and are currently being tested.

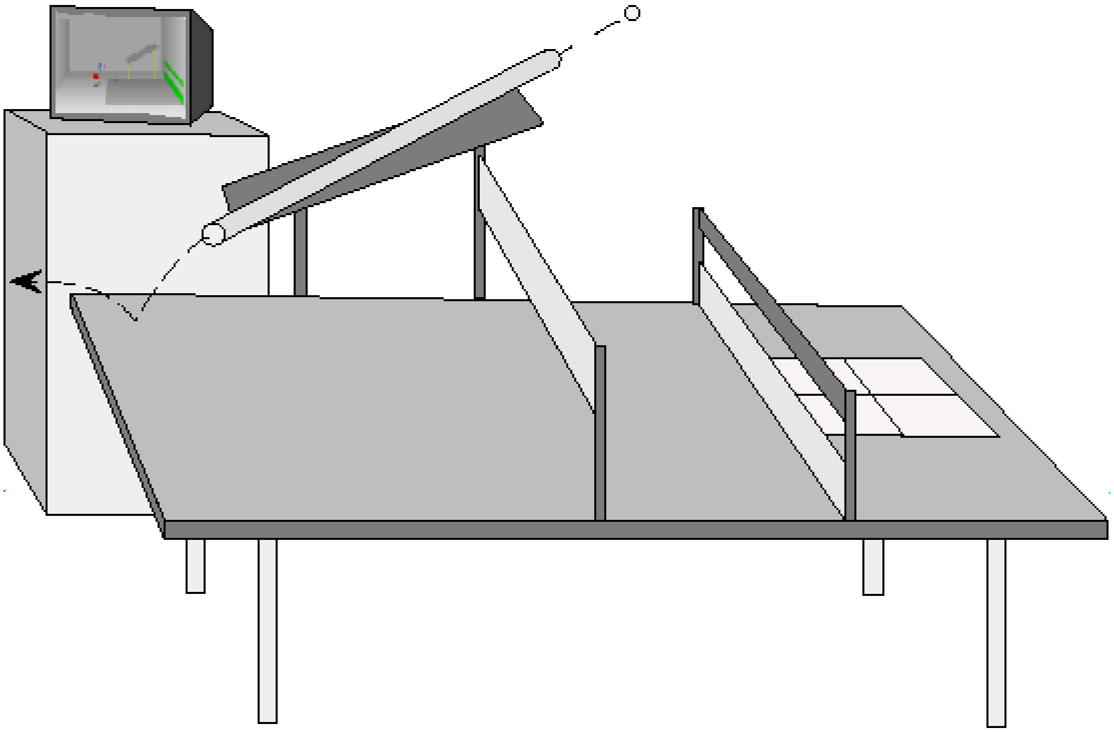

Figure 2. A schematic diagram of the ping-pong experimental setup. Sujects were standing in front of the table, hitting ping-pong balls with a paddle held in the left hand, and trying to send them to the target zone. The balls had to pass through the opening, and under the obstacle in the middle of the table. During training the subjects were standing in the same place required for the real task, and looking at the monitor.

In the two systems described below, both the desired and actual movements were displayed in real time in a virtual environment using a desktop computer. Our main concern with using purely virtual training was that transfer from virtual to real tasks has sometimes been difficult to achieve in prior studies. To avoid such failure, we made every effort to "integrate" the virtual environment in the physical context where the task normally takes place.

4.4.1 Scenes and Teacher. We created virtual worlds closely matching the desired task, i.e. containing the relevant objects with appropriate scale and distances among them. Instead of a fully immersive environment, we used a computer monitor that was properly aligned with the physical setting (as shown in Figure 2). In the ping-pong experiment described below, we interleaved training in the virtual environment with practice of the real task.

A human expert/therapist executed a desired movement, which was recorded with a motion sensor (Polhemus FASTRAK) and imported in the simulation. An animated object (the Teacher object) displayed the desired movement. Another animated object (the User object) displayed with minimum delay the position and orientation of a physical object held by the user. The user was instructed to repeatedly produce movements such that the User and Teacher objects remained perfectly aligned at all times. Additional examples of scenes are shown in Figures 3, 6, and 8. To aid the acquisition of the desired movement, we implemented a number of additional features, described below. The movement parser, score calculation, and semi-automated training features were developed based on our experience with the system in experiments 1 and 2. They are being evaluated in our current experiments.

4.4.2 Practice Modes. In the main practice mode the Teacher animation was displayed repeatedly, while the user attempted to synchronize with it and track the desired movement accurately. It was possible to speed up or slow down the prerecorded animation, trigger the animation when the user started moving (see Movement parser), or display a static 3D trace of the entire Teacher movement along with the animation (see Figure 6).

It was also possible to work in passive mode, in which the Teacher "followed" the User while still remaining on the prerecorded trajectory. This was accomplished by rewinding the prerecorded animation (at each point in time) to the videoframe in which the Teacher object was as close as possible to the current position and orientation of the User's object. This provided a way for the user to learn the spatial path of the movement separately from the velocity profile. We also included an active component in this mode, where the Teacher moved a small distance along its trajectory if the User made no movements for a certain period of time. This provided guidance especially helpful for patients.

4.4.3 Augmented Feedback. In both practice modes, the mismatch between User and Teacher movements could be emphasized by connecting the two objects with lines, or plotting circular arrows in the direction of the orientation error. The movements could be recorded and replayed later. Sound cues could be added to help timing. The visual display of either the teacher or patient trajectories could be altered so that the trajectory appeared as a wire frame or solid object. The size and frame frequency of teacher and patient ‘objects’ could also be varied.

4.4.4 Movement Parser. The repeated Teacher animation described above required the user to synchronize with the system – which was difficult for some patients. To alleviate this problem we developed an automated movement parser, which monitored the FASTRAK sensor data and determined online when the user started and stopped moving. This information was used to trigger the Teacher animation, and also to provide sound cues.

We have found it rather difficult to implement a movement parser that works reliably. After experimenting with various designs, we opted for a semi-automatic solution: the system provides a number of settings specifying velocity thresholds, distance thresholds, etc. and ways of combining them in Boolean expressions. These settings are adjusted by an expert user for each desired trajectory being trained with the system.

4.4.5 Score Calculation. An online scoring system has been developed to provide objective feedback as to how well the user is executing the desired movement. Since different movements require different scoring techniques, we implemented a general-purpose algorithm with a number of parameters that an expert user could set. The algorithm takes into account characteristics of the movement (average speed, duration, smoothness, etc.) as well as “closeness” to the desired trajectory. The latter is specified by entering weights for different types of error: position, orientation and velocity deviations, and different temporal misalignments.

4.4.6 Semi-automated training. It was clear from the beginning that a fully automated training schedule would be unrealistic. In both training and rehabilitation, the system was originally operated by an expert instructor/therapist who controlled the transitions between different scenes, training modes and feedback features. The expert was guided to a large extent by observing the subject/ patient performance. In our most recent version, however, we have implemented a semi-automated 'script' feature. The script feature allows us to specify prior to training, a series of scenes, teacher and feedback settings, that progress automatically through a preset number of repetitions, simulating a treatment or training session. The scores computed online can be linked to this function, so that the script will 'jump ahead' if the task is too easy for the subject. The level of difficulty can be adjusted by changing the score value set as the goal. In general, our intention was to utilize the valuable expertise of a human instructor/therapist as much as possible, while automating routine tasks and providing new sources of real-time feedback and motivation.

5. Experiment 1: Training a Complex Motor Skill in Normal Subjects

5.1 System Design

Our first system was designed for training normal subjects to hit ping-pong balls and send them into a target zone.(Todorov et al. 1997). The setup is shown in Figure 2.

Subjects were standing in front of the table, holding a ping-pong paddle in their left hand (to make the task more difficult). A ball was dropped through the tube shown in the figure, at an angle and speed corresponding to an intermediate to difficult shot. The task was to hit the ball after it bounced from the table, and send it in the target area through the opening above the net. In the second experiment an additional obstacle was added in the middle of the table, in which case the ball had to pass under it.

During training the subject looked at the monitor, which was properly aligned with the physical setup. The desired movement was recorded from an expert hitting a real ball on the real task. When the recorded expert movement was transferred in the virtual environment, the spatial alignment was preserved. In other words, in order to track the desired movement in the virtual environment, the subject had to make a movement in the real wold which would have accomplished the real task.

A snapshot of the simulation is shown in Figure 3. We used shadows, illumination cues and occlusions to create a realistic simulation. The red paddle shows the User position and orientation, the transparent blue paddle shows the Teacher recording.

The ball shown in the figure only interacted with the Teacher paddle; its dynamic simulation was adjusted so that the impact looked realistic. We did not allow the subjects to interact with the simulated ball during training, since our simulation of the ball dynamics was not completely accurate. One might wonder then, why include the ball in the first place if it does not seem to provide any additional information? It turned out however that including the ball was crucial - without it learning in the simulator did not transfer to the real task (see below).

During the development and fine-tuning of the system we experienced two main problems. Subjects occasionally confused the Teacher and User objects. This resulted in an intriguing but very undesirable

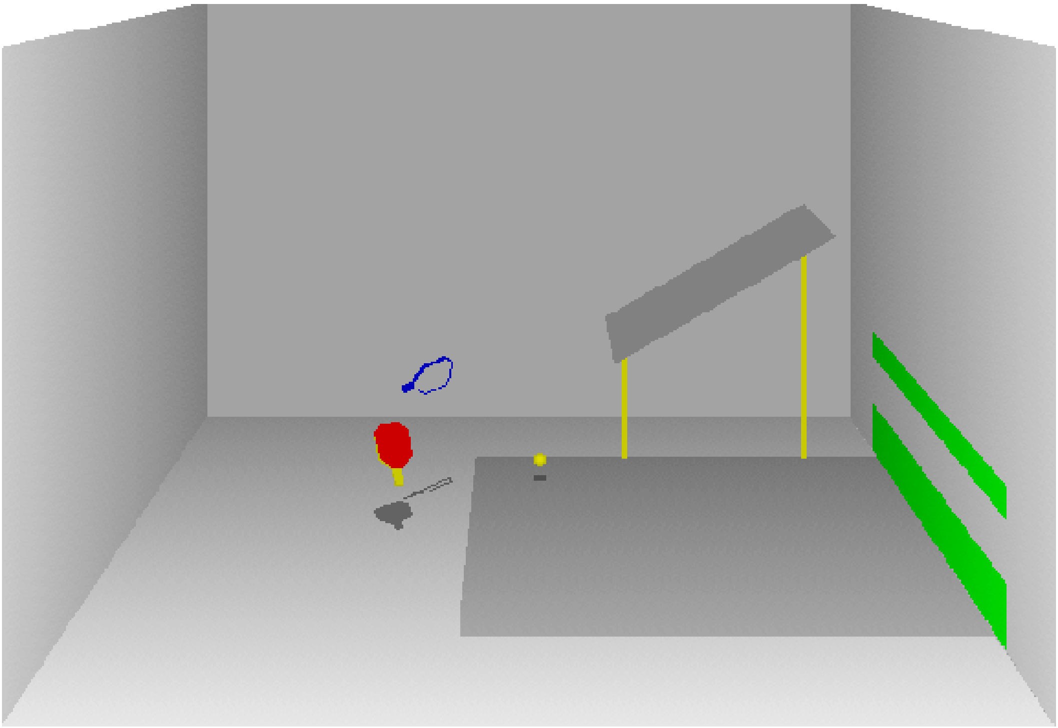

Figure 3. A snapshot of the simulation. The transparent paddle displayed a prerecorded (teacher) animation, which the user attempted to follow in real time with the red paddle. The simulated ball only interacted with the teacher paddle. The display was refreshed at about 40Hz.

phenomenon: when the User was falling behind the Teacher and instead thought he was ahead, the User movement slowed down further, causing an even bigger error. To avoid such confusion we had to make the two paddles as distinct as possible – the transparent Teacher paddle shown in Figure 3 solved the problem.

The other problem was synchronizing the two animations. We attempted to use a version of the movement parser described above, but that did not produce the desired effect. In order to make the parser robust (i.e. avoid signaling movement onset due to sensor noise), we had to use reasonably high velocity thresholds. As a result the system detected movement onset with a delay, so the Teacher animation started after the User movement. Such timing misalignment was very detrimental – the subjects spend most of their time trying to compensate for it rather than learning the teacher trajectory. We solved this problem by simply animating the Teacher at fixed intervals, and asking the subject to synchronize with it.

5.2 Experimental Results

In Experiment 1, volunteer college students were randomly assigned to one of three groups - a pilot group with 13 subjects, a control group with 20 subjects and a training group with 19 subjects. Each subject was introduced to the apparatus and the task, and given 10 practice balls, the score for which was not recorded. Then a baseline was recorded over 50 trials, which lasted approximately 10 min. Subjects in all groups

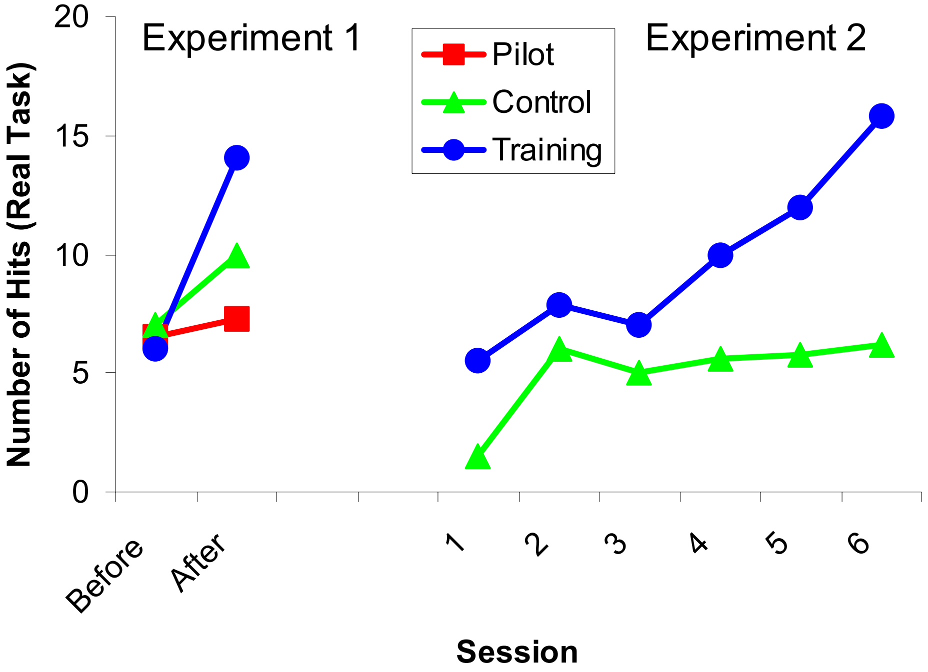

Figure 4. Summary of experimental results, in numbers of successful trials (in Experiment 1 the maximum was 50, in Experiment 2 the maximum was 60)

were able to see where the ball landed and therefore knew their score, but were not given any other feedback. After the first block the control group was given standard coaching by the experimenter, who was an experienced player. Coaching included discussion of what the subject was doing wrong, demonstration, and extra practice balls.

Subjects in the pilot and training groups were trained in the simulator, for the same amount of time as the coaching given to the control group. Training started with a slow version of the teacher trajectory, followed by 1-2 min of passive mode training. After that only the main mode (repeating the desired movement with the teacher) was used. The only difference between the pilot and training groups was that the pilot subjects were not shown a ball in the virtual environment, while the subjects in the training group could see the teacher paddle hitting a simulated ball.

Results are shown in Figure 4. Analysis of variance showed that all groups started with similar performance on the first 50 trials (no significant differences were found). At the end of the experiment the pilot group did not have better performance compared to the control group; pilot subjects were actually worse, and their improvement was smaller (the differences were close to being significant). The training group, however, had significantly higher scores on the second block of trials (p=0.02), and more improvement (p=0.01), compared to the control group.

In Experiment 2, we made the task more difficult by increasing the speed of the bouncing ball, and introducing the obstacle in the middle of the table (Figure 2). This time we included 6 sessions (2 sessions per day in 3 consecutive days). In each session, control subjects played 30 practice balls, followed by 60 test balls (scores in Figure 4). Training subjects used the simulator for the same amount of time that it took the controls to play the 30 practice balls. Then the training subjects were also tested on 60 balls in the real task. Again, we found significantly higher ($p=0.02$) performance in the training compared to the control group on the last session. The improvement of the training group was also significantly higher ($p=0.05$).

In summary, in two experiments involving more than 70 subjects we demonstrated that training in the virtual environment resulted in better performance than a comparable amount of coaching or extra practice on the real task.

It should be stressed however that while the performance gains we observed are rather encouraging, our analysis revealed a complex relationship between the teacher trajectory and success on the task. On one hand, the increased performance of the training group suggests that learning a teacher trajectory is useful. On the other hand, if accurate reproduction of the teacher trajectory was related to task performance, we would expect to see within the training group a correlation between each subject's performance level and how well that subject followed the teacher trajectory. No such correlation was found. The trajectories used on the real task deviated significantly from the teacher trajectory, and for most subjects these deviations increased over the three days (Figure 5) while performance was improving. A related result was found by Brisson and Alain (1996) in a simpler interception task. These authors showed that training to reproduce a teacher trajectory improved performance, but the details of the teacher trajectory did not make a difference. In other words, imitating a teacher trajectory in the context of the task helps, but not because the exact trajectory being imitated was learned. The exact reason for this improvement remains unclear.

Furthermore, only when the teacher trajectory was coupled with a simulation of a ball did the training subjects improve in the real task. This unexpected result suggests that seemingly minor changes in the simulator can dramatically reduce its efficacy, while the details of the teacher trajectory remain identical. One possible explanation for this effect is that the ball simplified the synchronization task - since its flight could have been used to predict when the Teacher paddle started moving.

5.3 Ongoing and Future Work

So far we have focused on training end-effector trajectories, i.e. only the paddle was displayed rather than the movement of the entire arm or body. While this is a good place to start, it certainly doesn't exhaust the possibilities. As Figure 1 suggests, this form of training can be used for more complex trajectories involving multiple body parts. For example, consider the game of golf. In that case there is a well defined end-effector (the golf club), however training the movement of the golf club alone is probably insufficient. Instead golf coaches insist that the movement of the entire body is important. Another good example of a task where whole body movement matters is dance – in that case there is no end-effector to begin with.

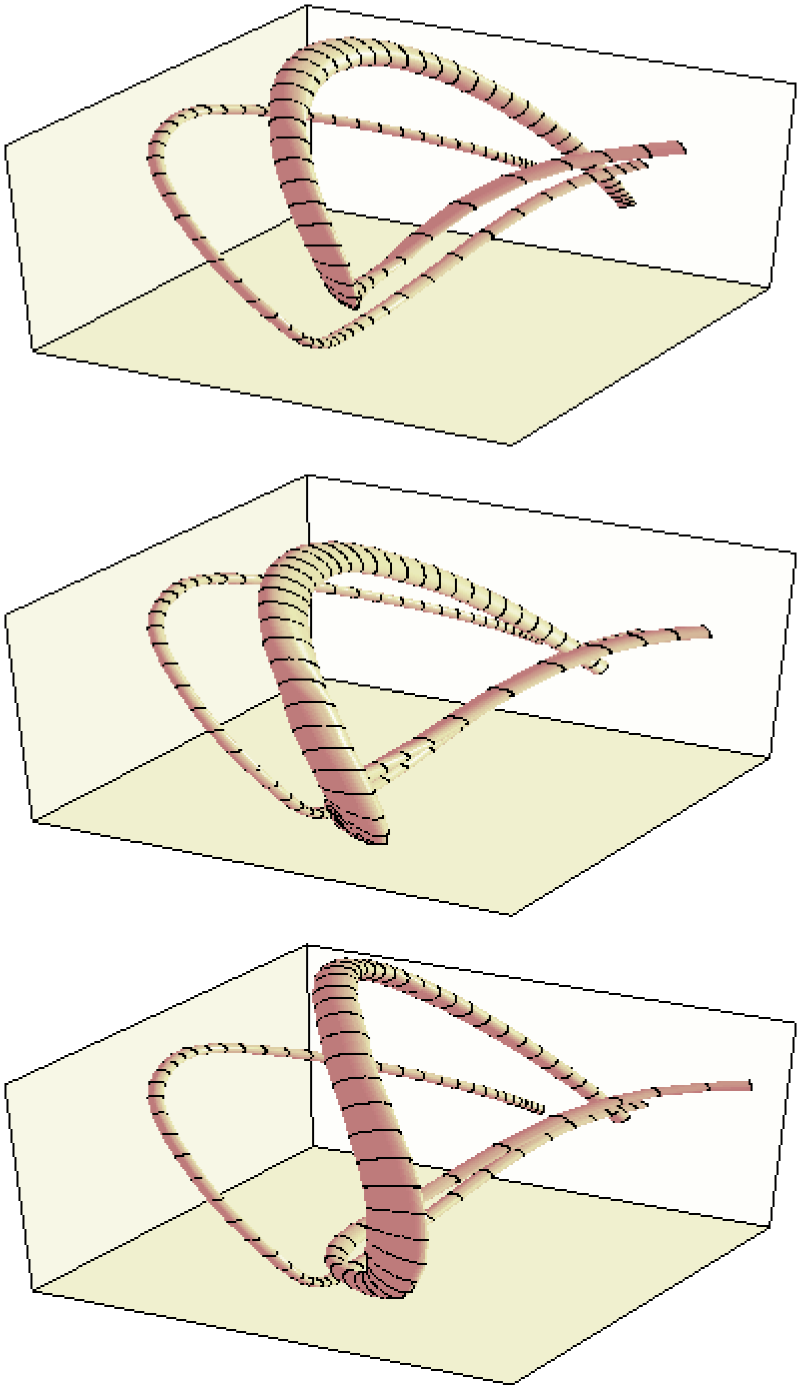

Figure 5. Trajectories of one training subject on the real task, days 1-3 from top to bottom. The thin curve is the teacher trajectory. The thick curve is the average subject's trajectory on that day, where the width corresponds to one standard error of the mean. Black rings mark equally spaced points in time, i.e. wider ring spacing corresponds to higher velocity. Note that the deviation between teacher and subject movements (on the real task) actually increases with practice, except for the region of impact with the ball (lower part).

If our training method is to be extended to movements of complex objects with multiple joints, the problem of confusing the Teacher and User objects becomes much more severe. Preliminary attempts to do that with the system shown in Figure 1 indicate that more advanced visualization techniques are needed before the motion of the entire hand can be tracked successfully. Once the User and Teacher hands are overlaid, there are so many small discrepancies that it is difficult to focus on any one of them and correct for it. One method we are considering is an automated procedure that selects the body parts with largest errors, and highlights them – by making all other parts semitransparent for example.

In the case of hand posture, we recently studied a matching task (Todorov and Ghahramani 2000) where subjects were shown computer generated images of random hand postures and asked to reproduced them as accurately as possible. Hand joint angles were recorded using a Cyberglove. The surprising finding was that subjects were very inaccurate in matching the individual joint angles: the average correlation coefficients (R^2) of desired and actual values ranged between 0.05 and 0.2. It turned out however that certain changes in overall hand shape (i.e. linear combinations of joints identified through Canonical Correlation Analysis) were matched better than any single joint: the first five such combinations had correlation coefficients of about 0.8! This result suggests that in the case of multi-articulate bodies, the human visual system may extract overall shape rather than individual joint angles. Information about such visually salient shape changes can be useful in the future for designing better feedback systems.

6. Experiment 2: Rehabilitation in Patients with Stroke

6.1 System Design

The second system was designed for the purpose of arm rehabilitation. The training method was very similar to the ping-pong study, however the rationale behind it was somewhat different. Consider a stroke patient who is currently incapable of lifting the impaired arm above shoulder level. What “desired movement” should be chosen for such a patient? There is an infinite variety of tasks (involving objects above shoulder level) which the patient is incapable of executing. This is where the skill of an experienced therapist is required. Based on the therapy evaluation, the therapist will determine the family of movements that are likely to be useful for a particular patient, and the family of tasks where such movements could arise naturally. In other words, we are now starting with a desired movement and constructing a task around it. This is the opposite to training a specific motor skill, where the task is primary and the desired movement is the one used by experts on that task.

Figure 6 illustrates the main features of the system. The trace corresponds to a desired movement, which the User (cube located near ‘hand’) is attempting to follow. The simplest way to create a task compatible with a given movement is to insert objects along the movement trajectory: in this case the donut was placed around the desired path, a second cube marked the beginning/end of the movement, and the task was defined as starting from the second cube and passing through the donut.

The major technical problem we experienced with the rehabilitation system was alignment of the patient with the virtual environment. A scene contains a number of virtual objects and a teacher recording,

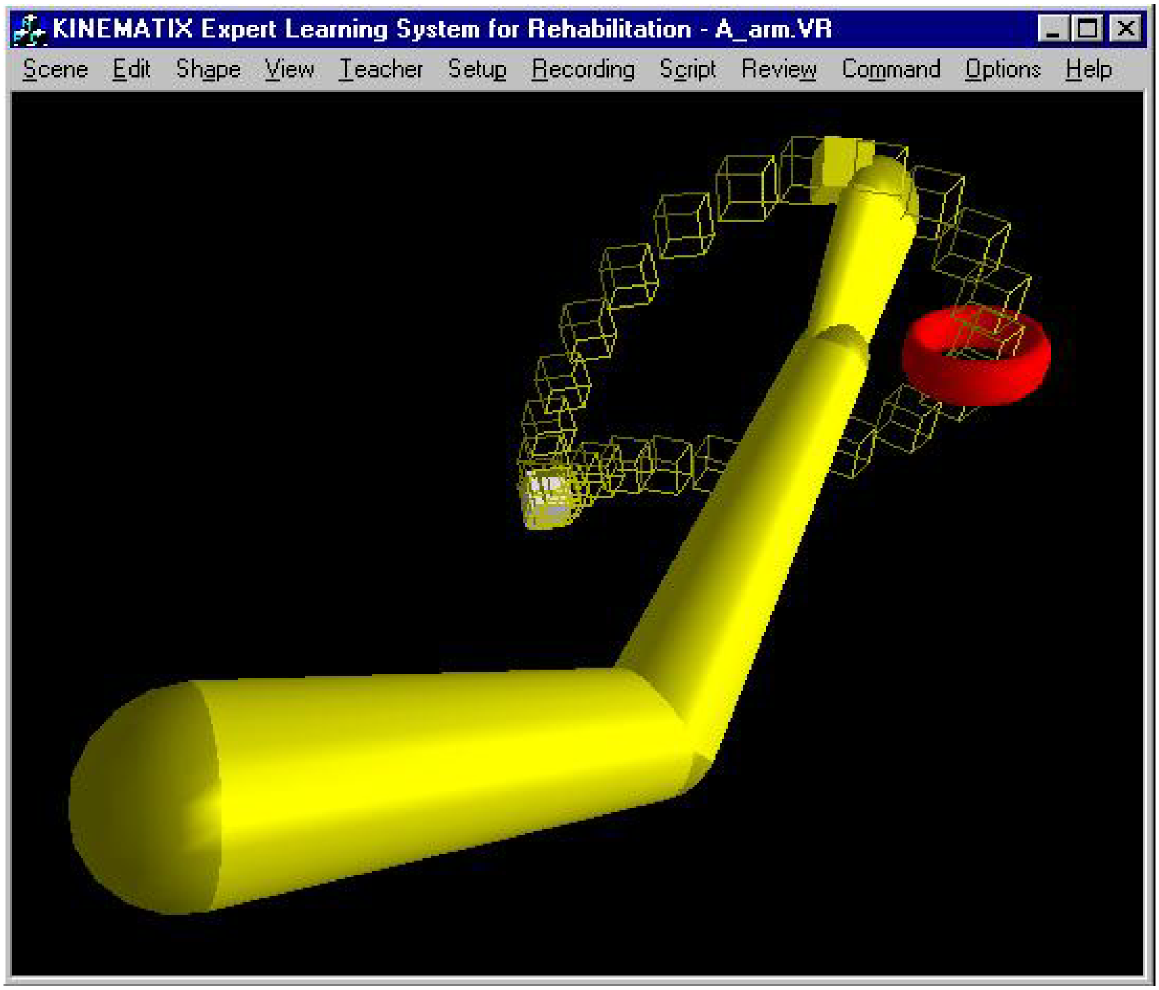

Figure 6. An example of a virtual scene in the rehabilitation system (not used for training). The arm of the patient was tracked in real time, the joint angles estimated, and the virtual arm rendered at about 40Hz. The trace shows the teacher tra jectory, for a task defined as starting from the white cube and passing through the donut.

whose positions and orientations are determined in a coordinate system centered at the transmitter of the FASTRAK (and oriented along its axes). The sensor readings are also defined in that coordinate system. Thus, if we insert a virtual object corresponding to some sensor location, save the scene, and position the sensor at the same location in physical space on the next day, it will appear on top of the virtual object, assuming the transmitter did not move. Similarly, if we adjust the scene so that the staring position and target correspond to the desired arm configurations of the patient, that correspondence will hold as long as the body of the patient is in exactly the same position and orientation relative to the transmitter.

In general, such alignment was not always easy to achieve. Although we attempted to position the patient in the same place relative to the computer in each training session, small differences inevitably occurred. And even small differences in patient position affected the alignment of the patient and teacher trajectories. While normal subjects could easily move to realign themselves, patients were slower and had much more difficulty with this task. Therefore, these alignment problems were resolved in software.

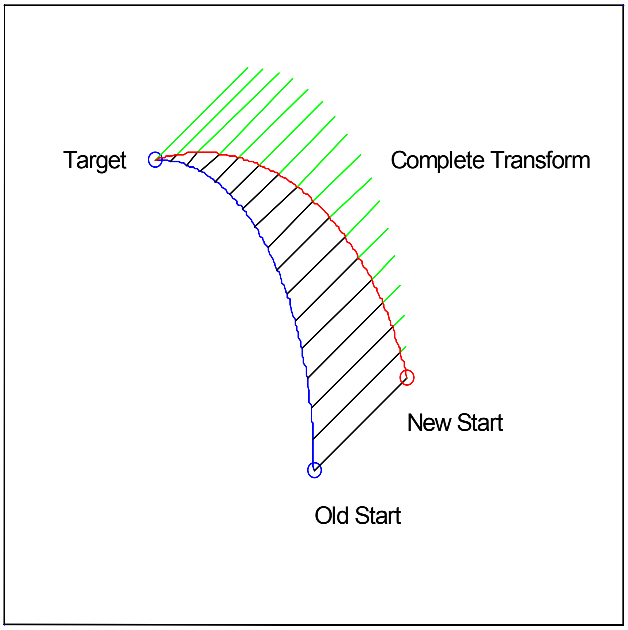

Figure 7. Illustration of the teacher alignment algorithm. The straight lines correspond to applying a transformation that takes the old starting position to the new starting position. To each frame of the teacher animation, we apply the same transormation scaled by the normalized distance (1-0) away from the starting position. This procedure results in a new teacher movement that looks very similar to the original, but starts at the new desired position.

6.1.1 Scene Alignment. One of the objects in the virtual scene (corresponding to the desired starting position and orientation of the movement) was defined as the "starting position object". To align the scene, we placed the physical sensor at the new starting position (and orientation). The system applied a rigid body transformation to the entire scene (i.e. a translation and a rotation) such that the old starting position object moved to the current sensor location. All other objects (and the teacher) moved accordingly, so that relative distances and orientations in the scene were preserved. Note that such a transformation is uniquely defined.

The problem with this procedure was that typically the whole scene tilted. In such cases we used a different alignment command that aligns only the positions of the physical sensor and the staring position object, but leaves the orientation of the scene intact.

6.1.2 Teacher Alignment. Sometimes the scene was aligned properly, but the patient was incapable of assuming the starting arm configuration of the teacher because of weakness, decreased range of joint motion, or muscle contractures. In such cases the system adapted the teacher movement, so that it started at a new starting position (one achievable by the patient) but ended in the old ending position, or passed though a previously specified target.

The algorithm used to solve this problem defined a transformation (which included both translation and rotation) that sent the beginning teacher frame into the new starting position. Also, the transformation could be scaled, to allow us to apply the complete transformation, half of it, etc. (for translations. scaling was straightforward; for rotations the quaternion representation was computed and the rotation angle was scaled; the rotation axis was unchanged). For each frame of the teacher animation, we determined how much of the transformation should be applied. For the beginning frame, the number was 1 (i.e. we want to transform the beginning frame completely to match the new starting position). For the target frame the number was 0 (i.e. we want to leave it unchanged). For intermediate frames we interpolated between 0 and 1. To preserve the shape of the speed profile of the teacher movement, the interpolation was based on distance from the target, rather than time. Figure 7 illustrates the procedure. The transformed teacher movements produced by this algorithm were surprisingly similar to the original.

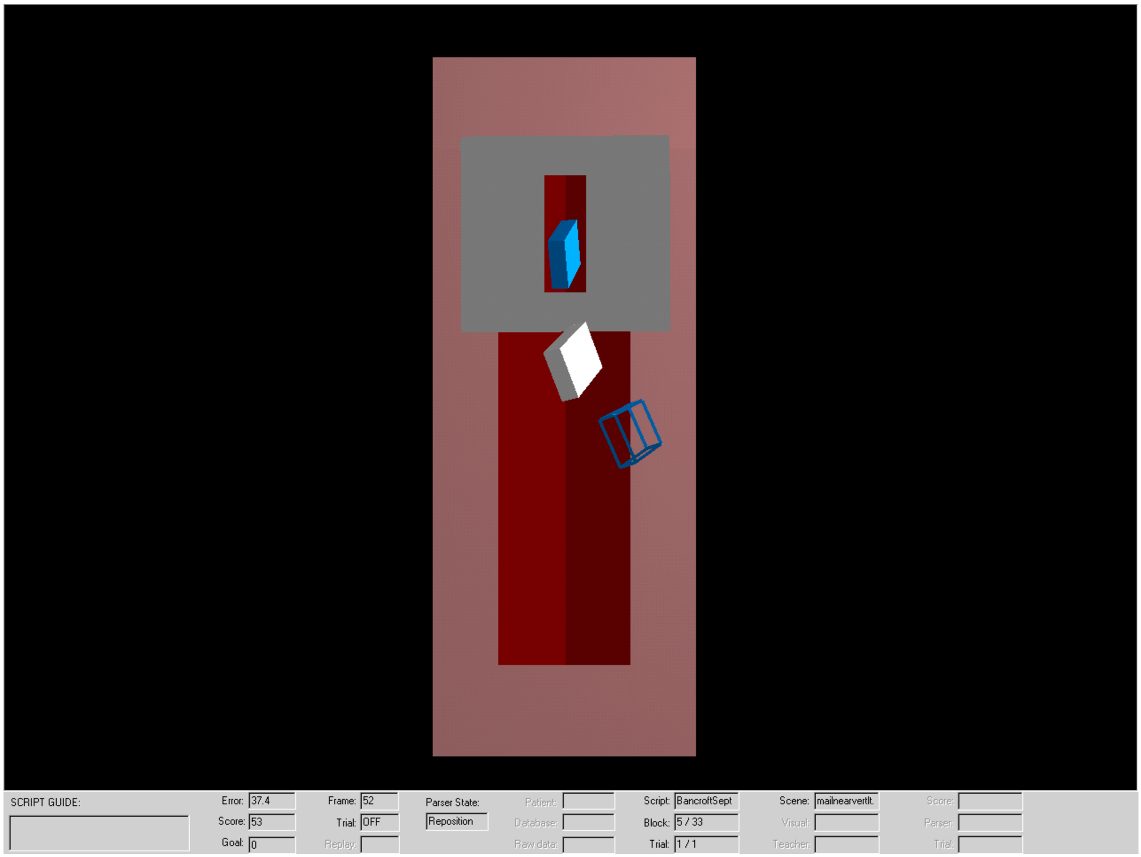

Error: 37.4
Score: 53
Goal: 0
Frame: 52
Trial: OFF
Parser State: Reposition
Script: BancroftSept
Block: 5 / 33
Trial: 1 / 1
Scene: mailnearvertlt.

Figure 8. Example of a ‘mailbox’ scene used for VE training. The ‘teacher’ envelope is shown entering the target slot. (In VE, the entire trajectory is animated, and objects appear in color against a black background.) Below the teacher, the ‘patient’ envelope is shown; this envelope moves when the patient moves his/her hand. Below, a wire-frame rectangle identifies the ‘start’ position, used to align the teacher and patient trajectories at movement onset

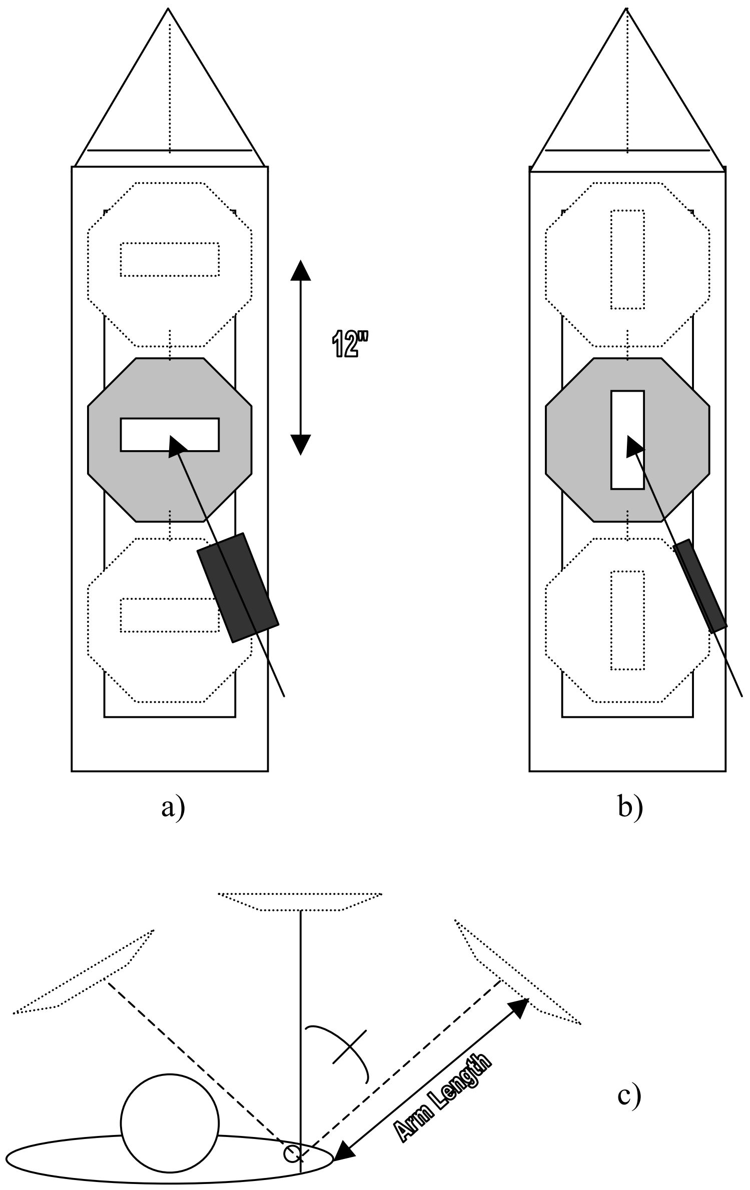

a) b)

c)

Figure 9. Schematic of real world reaching set-up used to test subjects before and after training in a similar virtual task. a) 3-D ‘mailbox’ with movable target slot (1”x5”); horizontal orientation was used for supinated and pronated reach attempts. ‘Envelope’ (4”x 8”x ¼”) is shown moving toward target. Center target was positioned at a location corresponding to hand position when UE was in 90^o shoulder flexion/ neutral abduction, and elbow fully extended (i.e., distance equal to arm length and height equal to height of acromion in sitting position). For testing upper and lower workspace reaching, the slot was positioned 12” above and below center position Target positions later trained in VE are shown in gray; untrained locations are shown in white. b) vertical slot orientation was used for reaching with forearm in neutral position; c) testing in transverse plane (view from above); slot was positioned 45^0 medial (adducted reach) and lateral (abducted reach) to the center position (shoulder-centered frame of reference). Only center forward position (gray) was trained in VE.

6.2 Experimental Results

In a pilot study (Holden et al., 1999) we treated two patients with hemiplegia using a computer-generated virtual environment (VE) to train upper extremity reach in the impaired limb. The goal of the experiment was to answer three questions: 1) Can hemiplegic subjects improve in a virtual task following virtual practice?; 2) Does learning which occurs in a virtual environment transfer to a similar real task?; and 3) Does learning in a virtual environment transfer to related but untrained real tasks, or to functional activities not specifically trained?

Subject 1was a 76 yr. old male, 3.5 yr. post left (L) cerebrovascular accident (CVA), due to thrombotic occlusion of the L internal carotid artery (ICA), confirmed by magnetic resonance imaging (MRI). He had resultant right (R) hemiparesis, significant expressive aphasia, but excellent receptive abilities. This stroke was his first, with no evidence of bilateral or brain stem stroke. He displayed no evidence of ongoing motor recovery in the right arm.

Subject 2 was a 76 yr. old female, 1.5 yr. post right (R) CVA, due to thrombotic occlusion of the right internal carotid artery, confirmed by computerized tomography (CT) scan. She had resultant left (L) hemiparesis. This stroke was her first, with no evidence of bilateral or brain stem stroke. She displayed no evidence of ongoing motor recovery in her left arm.

We selected only one movement for training in the virtual environment, with the criterion that it be a functional, goal-oriented movement that highlighted typical motor control problems seen in patients with stroke. To train this movement, we devised a reaching task, in which the subject held an ‘envelope’ (using a lateral grasp) then extended the arm to place the ‘envelope’ in a ‘mailbox’ slot.

Next we created a series of ‘scenes’ in the virtual environment. The scenes had a one to one spatial correspondence with the real world, and were displayed on a desktop computer. They were simple, containing only a virtual ‘mailbox’, and two virtual ‘envelopes’. One envelope was a ‘teacher’ who performed the correct movement over and over again (see Fig. 8). The teacher animation was a recording of a well-practiced normal subject performing the virtual task. The second envelope was a virtual representation of the real envelope that the patient held and moved during practice. Thus, the patient could match the endpoint trajectory of his/her movement with that of the teacher during training, in real time. (Movements of the patient were monitored using a Polhemus FASTRAK, then displayed on the computer in the context of the virtual scene.) The scenes progressed from easy to more difficult in order to train the movement in a sequential fashion. The sequence used was near and far reach, first with forearm pronation, then forearm neutral and lastly with supination. This resulted in three different hand orientations – palm down palm, facing right (when left arm was trained) or left (when right arm was trained), and palm up. The endpoint of the near reach for each different hand orientation was set at the distance midpoint of the normal trajectory for the far reach for that hand orientation.

The ‘teacher’ movement could be altered in a number of ways designed to facilitate learning. The animation could be adjusted in speed to any level slower or faster than the original, it could be paused at any point, displayed as a trace (solid line vs. moving object), or hidden from view entirely. The teacher

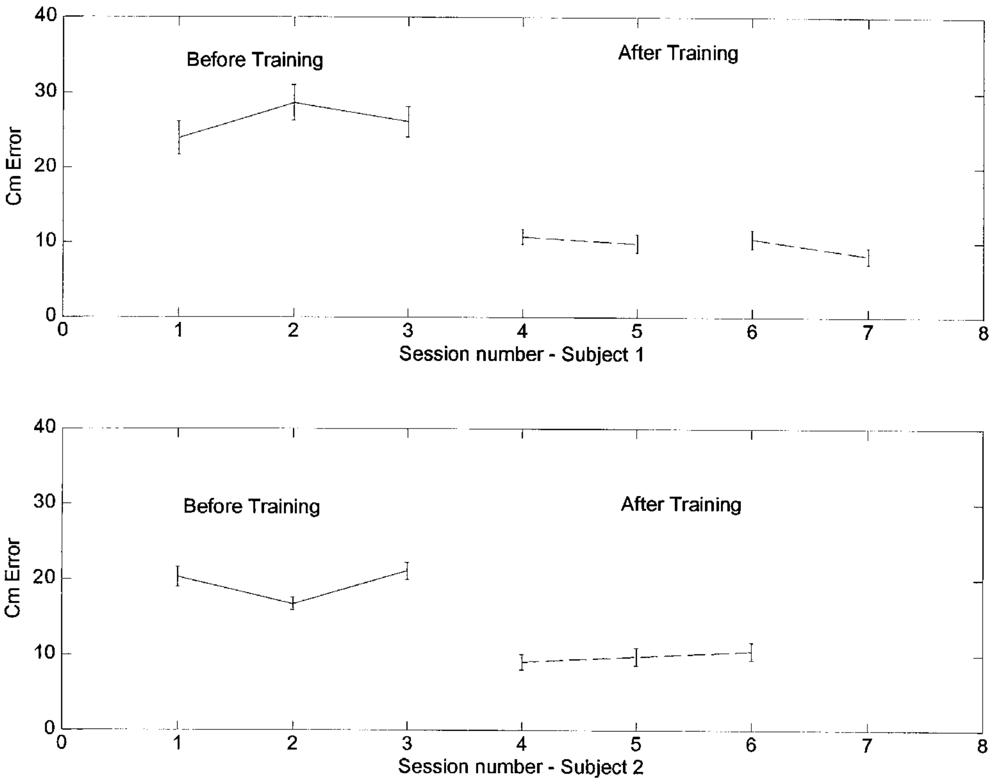

Figure 10. Performance on real-world reaching task pre/post virtual training. *Top Panel:* Mean error scores for distance (cm.) from target (averaged across all 9 target positions) plotted by session number for subject 1, involved arm. Sessions 1-3 (solid line) = pre-training scores; session 4-5 (dotted line) = scores after 8 training sessions; session 6-7 (dotted line) = scores after 16 training sessions. Magnitude of the pre/post differences would be considered statistically significant, based on non-overlap of the standard deviation values. Values for non-involved arm are not shown, but were ~ 1 cm with little variability for all sessions, with the exception of session 1, when error value was ~ 5 cm with +/- 2 cm standard deviation. *Bottom Panel:* Mean error scores for distance (cm.) from target (averaged across all 9 target positions) plotted by session number for subject 2, involved arm. Sessions 1-3 (solid line) = pre-training scores; session 4-6 (dotted line) = scores after 16 training sessions. Values for non-involved arm are not shown, but were ~ 1 cm with little variability for all sessions. (The values for the non-involved arm do not represent true 'errors', i.e., subjects could readily perform the task; rather they result from the sensor position on the envelope.)

could also be made to 'follow' the trajectory of the subject, i.e., temporal information could be eliminated, and the learner could focus exclusively on learning the spatial elements of the movement. Additional features (audio and visual) were designed to assist the patient with error detection, timing and positioning. A model of the entire arm, as well as the held object, could also be displayed on the screen in real time, if desired. This was used intermittently in training, to correct the patient's tendency to try to achieve the task with compensatory movement patterns such as excessive shoulder abduction with elbow flexion, which were ineffective.

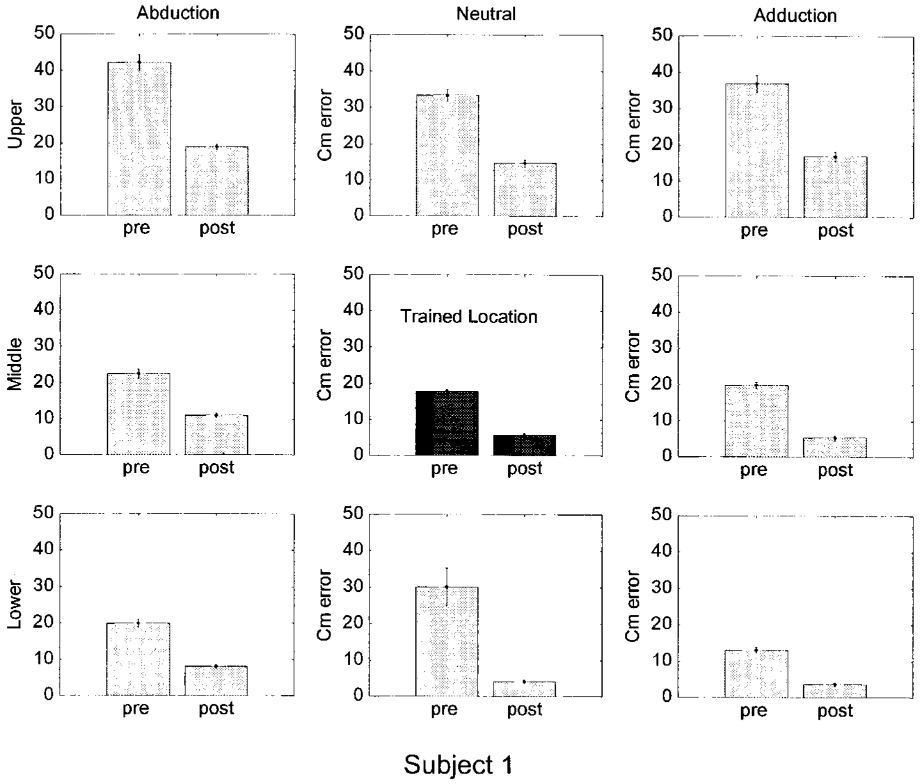

Figure 11. Distance errors (cm.) for real-world reaching task, plotted by target location in the workspace, for subject 1, involved arm. Values are means across three sessions (pre) and four sessions (post) training on a similar virtual task. Error scores for the trained movement are in the center of the middle row. Note that all locations show a reduction in error, despite no specific training in those locations. Note better transfer to upper parts of the workspace for this subject. Values for non-involved arm (not shown) were close to 1 cm and were nearly identical pre/post, with the exception of upper-adduction, which had a 10 cm mean error pre training, 1 cm post. (The 1 cm error score for non-involved arm was a function of the sensor position on the 'mail' piece, and not due to inability of the subject to perform the task.

Each subject received a series of 16 treatment sessions, of 1-2 hr duration, conducted by a physical therapist. The planned treatment frequency was 2x/ week, although subjects actually came less frequently than this due to problems with transportation and minor illnesses. Virtual performance improved over the 16 sessions, as measured by assessing how many scenes a subject could progress through in a session. (Once subjects could perform the movement in a scene correctly for three consecutive trials, they progressed to the next scene.) Real world performance in the 'mailbox' task showed improved reaching in both the trained and untrained parts of the workspace for both subjects. The distance and orientation error for the 'envelope' relative to the slot was calculated. (See Figure 9 for testing set-up)

S1 showed an average pre/post decrease in reach excursion error of 18 cm. (i.e., 64% reduction in error); S2 showed a decrease of 9 cm, representing a 50% reduction in error (Figure 10). The 9 and 18 cm gains represented roughly a 25% improvement in reach excursion. In both subjects, some of the largest improvements occurred in *untrained* parts of the workspace. S1 showed better transfer to upper regions of

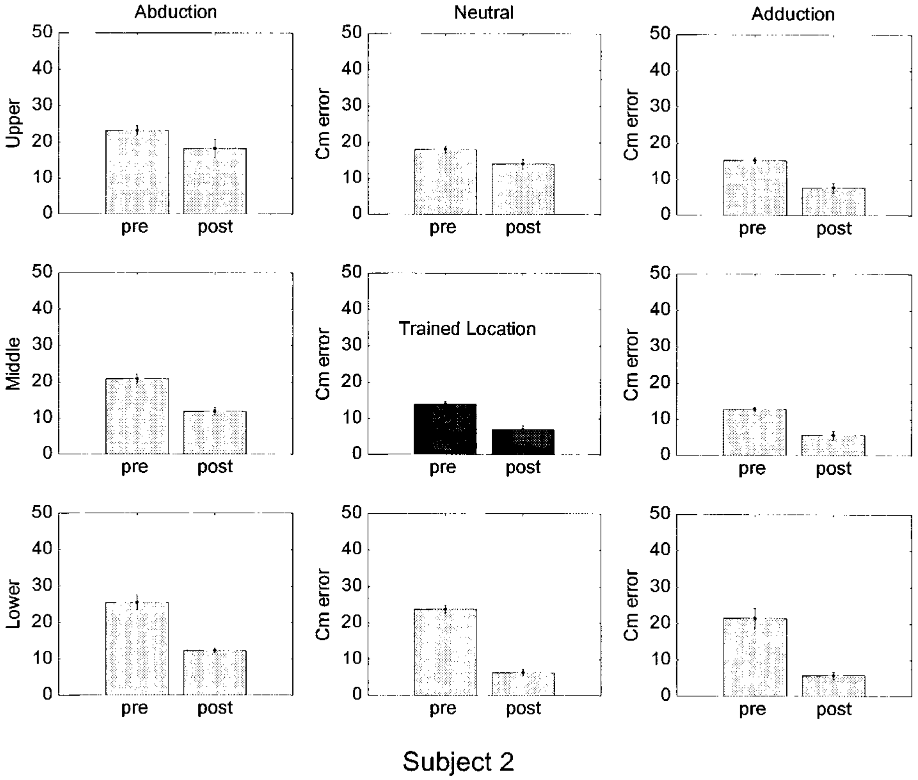

Figure 12. Distance errors (cm) for real-world reaching task, plotted by target location in the workspace, for subject 2, involved arm. Values are means across three sessions pre and post training on a similar virtual task. The trained movement is in the center of the middle row. However, most locations show a reduction in error, despite no specific training in those locations. Note better transfer to lower parts of the workspace for this subject. Non-involved arm values are not shown, but were ~1cm for all locations and essentially identical for pre/post tests.

the workspace; S2 showed better transfer to lower regions (Figures 11-12). Similar, but smaller magnitude results were found for hand orientation errors.

Several findings in this study have important clinical implications. The most important finding was that subjects not only improved on the virtual task, but also showed *transfer* of that improvement to similar real world tasks – both trained and untrained.

6.3 Ongoing and Future Work

The VE training system we have developed is currently being utilized in Italy in a clinical study of stroke patients who have upper extremity impairments (Piron, Dam, Trivello et al., 1999; Piron, Trivello, Cenni, et al., 2000). This study is an ongoing randomized, controlled clinical trial designed to evaluate the effectiveness of virtual environment training as a complementary therapy for motor rehabilitation following stroke. Experimental group subjects receive 1 hour of virtual environment therapy plus 1 hour of conventional physical therapy, daily for 4 weeks (20 two-hour treatments total); control group subjects are treated for 2 hours of conventional therapy daily for 4 weeks (20 two-hour treatments total). Both acute (< 6 months following stroke) and chronic (6-72 months following stroke) duration patients are being studied. Therapy in both virtual and conventional sessions is focused on the arm. All patients are evaluated before and after training, using the Fugl-Meyer scale for motor recovery (Fugl-Meyer, Jaasko, Leyman, et al., 1975) and the Functional Independence Measure (FIM) (Hamilton, Granger, Sherwin, et al., 1987), a scale that measures the amount of assistance needed to perform activities of daily living (such as eating and dressing). Fourty-Three subjects (28 experimental and 15 control) have been tested to date. Preliminary analysis suggests that the subjects who received the combined virtual + conventional therapy improved more than subjects who received conventional therapy alone, on both the test of motor recovery (Fugl-Meyer arm sub-score) and the functional scale (FIM). However, the study is still ongoing, and detailed statistical testing has not been performed. The response to the virtual environment training on the part of both the therapists and the patients who have used the system has been quite positive. No negative side effects have been reported, and both patients and therapist feel it provides enhanced motivation for the rehabilitation program.

We have recently begun a second study utilizing the VE system to study aspects motor learning and motor generalization in patients with stroke. In addition to knowledge about generalization, we hope to gain further practical experience with the newly developed features of our system (scoring, movement parser, and semi-automated functions.)

An extension of this work, currently being developed, is a networked virtual environment, which will allow the therapist to work with patients remotely. The current software is being modified so that new virtual scenes can be loaded remotely, alignment commands can be applied, and key feedback features can be turned on/off. We are also incorporating an internet videophone to provide an audio-visual link between the patient and therapist.

A third ongoing project is a study which uses the VE system to train subjects with acquired brain injury (ABI) to learn a task-oriented movement with the upper extremity. The study is being conducted at Bancroft NeruoHealth in Haddonfield, NJ. The task we are presently working with is pouring from a cup held in the impaired hand, at different locations in the workspace. During training, subjects hold a real cup in their hand while practicing the pouring movement and viewing their virtual performance (and that of the teacher) on the computer screen.

One purpose of this study is to gain more knowledge about how different clinical symptoms and impairments affect the ability of subjects with ABI to benefit from training with the system, and to identify equipment and software adaptations might enhance the ability of subjects with acquired brain injuries to improve from motor training in VE.

In the future, we would like to make the virtual environment used in our system more interactive and also incorporate a haptic interface. This would allow training of a greater variety of movements in more interesting ways. Currently static objects define the task, and the only moving parts are the teacher and patient animations. One can imagine instead a game-like scenario, where the objects respond to patient movements in various ways. For many patients, the response would need to be very simple, but for others we could adapt state of the art computer games with advanced graphics, sound effects, realistic physics engines and entertaining story lines to their levels. The problem is that current computer games are designed with traditional input devices in mind, such as a keyboard or joystick or mouse. For the purpose of motor rehabilitation we need games that accept input from 3D motion capture hardware, and use it in some meaningful way. If such games become available in the future they could be used not only for entertainment, but also as therapeutic tools.

7. Summary

In this chapter, we first reviewed the literature relevant to spatial-motor learning in VE in normal subjects, as well as the literature on clinical applications of VE in rehabilitation. Next, we described a system that we developed to enhance motor learning in healthy and disabled individuals. The system is based on learning by imitation and uses a virtual environment to implement training routines. We ended with a description of our experimental results to date, and a discussion of design issues that arose during our use of the system in all these experiments., along with a brief description of ongoing work and future research directions.

References

Atkeson, C. & Hollerbach, J. (1985). Kinematic features of unrestrained vertical arm movements. Journal of Neuroscience 5, 2318-2330.

Biocca, F.A. & Rolland, J.P. (1998). Virtual eyes can rearrange your body: Adaptation to virtual eye location in see-through head-mounted displays. Presence, 7(3), 262-277.

Bower, G. H. (1972). A selective review of organizational factors in memory. In E. Tulving & W. Donaldson (Eds.), Organization of memory (pp.93-137). New York: Academic Press.

Brisson T. and Alain C. (1996). Should common optimal movement patterns be identified as the criterion to be achieved? Journal of Motor Behavior, 28(3), 211.

Christansen, C., Abreu, B., Ottenbacher, K., Huffman, K., Masel, B., & Culpepper, R. (1998). Task performance in virtual environments used for cognitive rehabilitation after traumatic brain injury. Arch. Phys. Med. Rehabil., 79, 888-892.

Cobb, S.V.G. (1999). Measurement of postural stability before and after immersion in a virtual environment. Applied Ergonimics, 30, 47-57.

Cunningham, D. (1999). Therapy for New Millennium. Rehab Management, Dec/Jan, 84-85.

Darken, R.R., Allard, T., & Achille, L. B. (1998). Spatial orientation and wayfinding in large-scale virtual spaces: An introduction. Presence, 7(2), 101-107.

Docimo, S.G., Moore, R.G., & Kavoussi, L.R., (1997). Telerobotic surgery is clinical reality: Current experience with telementoring in adults and children. Presence, 6(2), 173-178.

Emmett A. (1994). Virtual reality helps steady the gait of Parkinson's patients. Computer Graphics World, 17, 17-18.

Foreman, N., Orencas, C. Nicholas, E., Morton, P., & Gell, M. (1989). Spatial awareness in seven to 11-year old physically handicapped children in mainstream schools. European Journal of Special Needs Education, 4, 171-178.

Fugl-Meyer, A.R., Jaasko, L., Leyman, I., Olsson, S., & Steglind, S. (1975) The post-stroke hemiplegic patient: A method for evaluation of physical performance. Scandanavian Journal of Rehabilitation Medicine 7,13-31.

Gillner, S., & Mallot, H.A. (1998). Navigation and acquisition of spatial knowledge in a virtual maze. Journal of Cognitive Neuroscience, 10(4), 445-463.

Greenleaf, W. & Tovar, M.A. (1993). Augmenting reality in rehabilitation medicine. Artificial Intelligence in Medicine, 6, 289-299.

Hamilton, B.B., Granger, C.V. Sherwin, F.S., Zielezny, M., & Tashman, J.S. (1987) A uniform national data system for medical rehabilitation. In M.J. Fuher (Ed.). Rehabilitation outcomes: Analysis and measurement. Baltimore: Brookes.

Holden, M., Todorov, E., Callahan, J., & Bizzi, E. (1999). Virtual environment training improves motor performance in two patients with stroke. Neurology Report, 23(2), 57-67.

Inman, D., Peaks, J., Loge, K., & Chen, V. (1994). Teaching orthopedically impaired children to drive motorized wheelchairs in virtual reality. Proceedings of the 2nd Annual International Conference on Virtual Reality and Persons with Disabilities. San Francisco, CA.

Kennedy, R.S., Lanham, D.S., Massey, C.J., Drexler, J.M., & Lilenthal, M. G. (1995): Gender differences in simulator sickness incidence; implications for military virtual reality systems. SAFE J. 25, 69-76.

Kennedy, R. S., Lilienthal, M. G., Berbaum, K. S., Baltzley, D.R., & McCauley, M.E. (1989). Simulator sickness in U.S. Navy flight simulators. Aviation, Space and Environmental Medicine, 60, 10-16.

Kozak, J.J., Hancock, P.A., Arthur, E.J., & Chrysler, S.T. (1993). Transfer of training from virtual reality. Ergonomics, 36, 777-784.

Kramer, J. & Leifer, L. (1989). The talking glove: A speaking aid for nonvocal deaf and deaf-blind individuals. RESNA, 12th Annual Conf., New Orleans.

Kuhlen, T. & Dohle, C. (1995). Virtual reality for physically disabled people. Comput. Biol. Med, 25(2), 205-211.

Lacquaniti, F., Terzuolo, C. & Viviani, P. (1983). The law relating the kinematic and figural aspects of drawing movement. Acta Psychol 54, 115-130.

Lackner, J.R. & DiZio, P. (1994). Rapid adaptation to Coriolis force perturbations of arm trajectory. Journal of Neurophysiology, 72, 299-313.

Lawson, B. D., Rupert, A.H., Guedry, F.E., Grissett, J. D., & Mead, A. M. (1997). The human-machine interface challenges of using virtual environment (VE) displays aboard centrifuge devices. In Proceedings of the Seventh International Conference on Human-Computer Interaction, (pp.945-948, Vol. 2). San Francisco, CA. Amsterdam: Elsevier-Science.

Lusted, H.S. & Knapp, B.R.. (1992). Biocontrollers for the physically disabled: a direct link from the nervous system to computer. Proc. Virtual Reality and Persons with Disabilities, CSUN, Los Angeles.

Morasso, P. (1981). Spatial control of arm movements. Experimental Brain Research 42, 223-227.

Nichols, S. (1999). Physical ergonomics of virtual environment use. Applied Ergonomics, 30, 79-90.

Piron, L., Dam, M., Trivello, E., Iaia, E., & Tonin, P. (1999). Virtual environment training ameliorates motor deficits in post-stroke patients. American Academy of Neurology Annual Conference, Toronto, Ontario, CANADA.

Piron, L., Trivello, E., Cenni, F., Iaia, E., Dam, M., & Tonin, P. (2000). Motor training in virtual environment for post-stroke and traumatic brain injured patients. 8th Annual Medicine Meets Virtual Reality Conference, January 27-30, Newport Beach, California.

Pugnette, L., Mendozzi, L, Motta, A., Cattaneo, A., Barbieri, E. & Brancotti, A. (1995). Evaluation and retraining of adults for cognitive impairments: Which role for virtual reality technology? Comput. Biol. Med., 25(2), 213-227.

Regian, J.W., & Shebilske, W.L. (1992). Virtual reality: An instructional medium for visual-spatial tasks. Journal of Communication, 42, 136-149.

Richard, P., Birebent, G., Coiffet, P., Burdea, G., Gomez, D., & Langrana, N. (1996). Effect of frame rate and force feedback on virtual object manipulation. Presence, 5 (1), 95-108.

Reiss, T. & Weghorst, S. (1995). Augmented reality in the treatment of Parkinson's disease. In Morgan, K., Satava, R.M., Sieburg, H.B., Mattheus, R., Christensen, J.P. (Eds.), Interactive technology and the new paradigm for healthcare (pp.415-422). Amsterdam: IOS Press.

Rolland, J.P., Biocca, F.A., Barlow, T., & Kancherla, A. (1995). Quantification of adaptation to virtual-eye location in see-through head-mounted displays. Virtual Reality Annual Int. Symp. '95, (pp. 55-66). Los Alamitos, CA: IEEE Computer Society Press.

Rose, F.D., Attree, E.A., Johnson, D.A. (1996). Virtual reality: An assistive technology in neurological rehabilitation. Current Opinion in Neurology, 9, 461-467.

Ruddle, R.A., Payne, S.J., Jones, D.M., (1999). Navigating large-scale virtual environments. Presence, 8, 57-68.

Schmidt, R. (1975). A schema theory of discrete motor skill learning. Psychological Review, 82, 225-260.

Schmidt, R. & Young, D. (1991). Methodology for motor learning: a paradigm for kinematic feedback. Journal of Motor Behavior 23, 13-24.

Shadmehr,R. & Mussa-Ivaldi, S. (1994). Adaptive representation of dynamics during learning of a motor task. Journal of Neuroscience 14, 3208-3224.

Smythe P. (1993). Use of networked VR to assist people with special needs. In Proc. First Ann Interntl Conf. Virtual Reality and Persons with Disabilities, (pp. 113-117). H.J Murphy.(Ed.), Northridge, California State University.

Stanney, K.M., Kennedy, R.S., Drexler, J.M., & Harm, D.L. (1999). Motion sickness and proprioceptive after effects following virtual environment exposure. Applied Ergonomics, 30, 27-38.

Stanney, K. & Salvendy, G. (1998). Aftereffects and sense of presence in virtual environments: Formulation of a research and development agenda. International Journal of Human-Comput. Interaction, 10, 135-187.

Theasby, P.J. (1992). The virtues of virtual reality. GEC Review, 7(3), 131-145.

Todorov, E. & Ghahramani Z. (2000). Degrees of freedom and hand synergies in manipulation tasks. Neural Control of Movement Abstracts, Key West, FL.

Todorov, E.., Shadmehr, R.., & Bizzi, E. (1997). Augmented feedback presented in a virtual environment accelerates learning of a difficult motor task. Journal of Motor Behavior, 29(2), 147-158.

Vamplew P. & Adams, P. (1992). The SLARTI system: applying artificial neural networks to sign language recognition. Proc. Virtual Reality and Persons with Disabilities, CSUN, Los Angeles.

Warner, D., Anderson, T., & Johanson, J. (1994). Bio-cybernetics – a biologically responsive interactive interface. Proc. Medicine Meets Virtual Reality, San Diego, CA.

Wilson, P.N., Foreman, N., & Stanton, D. (1997). Virtual reality, disability and rehabilitation. Disability and Rehabilitation, 19(6), 213-220.

Wilson, P.N., Foreman, N., & Tlauka, M. (1996). Transfer of spatial information from a virtual to a real environment in physically disabled children. Disability and Rehabilitation, 18, 663-637.

Wolpert, D., Ghahramani, Z. and Jordan, M. (1995). Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. Experimental Brain Research 103, 460-470.