

Training for Physical Tasks in Virtual Environments: Tai Chi

Philo Tan Chua
David Ventura

Rebecca Crivella
Todd Camill

Bo Daly
Jessica Hodgins
Entertainment Technology Center
Carnegie Mellon University

Russ Schaaf
Randy Pausch

Abstract

We present a wireless virtual reality system and a prototype full body Tai Chi training application. Our primary contribution is the creation of a virtual reality system that tracks the full body in a working volume of 4 meters by 5 meters by 2.3 meters high to produce an animated representation of the user with 42 degrees of freedom. This – combined with a lightweight (< 3 pounds) belt-worn video receiver and head-mounted display – provides a wide area, untethered virtual environment that allows exploration of new application areas. Our secondary contribution is our attempt to show that user interface techniques made possible by such a system can improve training for a full body motor task. We tested several immersive techniques, such as providing multiple copies of a teacher's body positioned around the student and allowing the student to superimpose his body directly over the virtual teacher. None of these techniques proved significantly better than mimicking traditional Tai Chi instruction, where we provided one virtual teacher directly in front of the student. We consider the implications of these findings for future motion training tasks.

1 Introduction

With recent advancements in wireless technology and motion tracking, virtual environments have the potential to deliver on some of their early promises, creating effective training environments and compelling entertainment experiences. In this paper, we describe and evaluate a system for training in a physical task – Tai Chi – and show the application of training techniques not possible with real world training.

The capture and display system is wireless, freeing the student from the encumbrance of trailing wires. We place a Tai Chi student in a 3D rendered environment with a virtual instructor, and they learn Tai Chi by mimicking expert motions, similar to real-world Tai Chi instruction. Our hypothesis is that virtual environments allow us to take advan-



Figure 1. Left: A student in the Tai Chi trainer. Right: The virtual world. Student is in light, teacher in dark clothes

tage of how physical activities such as Tai Chi are currently learned while visually enhancing the learning environment in ways not possible in the real world. The student's movement is recorded with an optical motion capture system that allows real-time capture of the whole body. Captured motion is used to animate representations of the student and teacher and displayed via a wireless head-mounted display (HMD), allowing students to have a natural representation of their body in the virtual environment and easily acquire feedback on their movement.

We view this system as a first step toward more effective and general systems for the training of physical tasks. Physical training is an economically important application with a significant impact on both health and safety. VR has great potential for enhancing physical training, with the ability to record a motion once and then replay and practice the motion unlimited times, enabling self-guided training and evaluation. Tai Chi is itself important because there are demonstrated positive effects in preventing falls in the elderly and other health improvements [22].

We chose Tai Chi because the slow nature of the movements and the range of motion fit the properties of our VR system well, additionally we can offer evaluation and feedback to the student during training. Tai Chi is a challenging training application because the sequence of movements, the *form*, is complicated and the standards for performance are exacting. Novices struggle to memorize the sequence of movements and more experienced students spend 20 years or more refining their form. Because emphasis is on balance and shape of the body during slow movements, students are able to adjust their motion based on feedback from the teacher during the sequence.

Figure 1 shows a user wearing the wireless virtual reality equipment (a belt and an HMD) and the virtual environment with the real-time data applied to a student model. In the virtual environment, the student is rendered with a white shirt and the master wears red.

Real-time, full body capture of the student's motion allows us to experiment with different forms of visual feedback. We implemented several user interaction techniques, such as providing multiple simultaneous copies of the teacher surrounding the student, and allowing the student to superimpose his body directly inside the teacher's body. While the former could be done via projection or CAVE-like systems, the latter technique is not readily implemented without the use of an HMD. We compared various interfaces objectively by computing the error in limb position between student and teacher, and subjectively via a questionnaire of our subjects. Surprisingly, both the objective measures and subjective questionnaires showed that none of the virtual environment enabling techniques performed better than merely mimicking a real world situation, with one teacher in front of the student.

2 Background

Wireless virtual environments with full body capture have become possible only in the past few years. To our knowledge, ours is the first wireless full body virtual environment system usable for a broad variety of applications. The only other wireless system of which we are aware was a highly specialized dismounted soldier training system that required the participant to wear 25 pounds of equipment [11], severely limiting its applicability. That system was constructed for combat training and used a General Reality CyberEye HMD and Premier Wireless for the transmission of a bi-ocular NTSC video image. The motion of the soldier and rifle were captured with a Datacube optical system.

The technology for wireless whole body capture in real-time has been available for a number of years, first with magnetic sensors and more recently with commercial optical motion capture systems. Researchers have explored

ways to make magnetic systems less burdensome for the user by using fewer sensors and supplementing the data with inverse kinematics for the limbs where motion is not directly measured [2, 15, 10]. The *HiBall* wireless, wide-area system created by Welch et al. provides very good data but requires the student to wear relatively encumbering trackers [21]. Technologies for capturing some of the motion of the user via video have also been developed [23].

With the system described in this paper, our goal has been specifically to explore the use of virtual environments in training applications. A number of virtual environments have been developed in which the goal is to teach the student decision-making strategies for such situations as hostage negotiations, earthquakes, parachute descent and fire fighting [13, 17, 6]. Fewer virtual environments have been built that attempt to train the user for physical tasks. Emering et al. created a system for action recognition that was applied to a martial arts scenario [1]. Davis and Blumberg built a virtual aerobics trainer in which the user's actions were captured and recognized as he or she was instructed to "get moving" or given feedback such as "good job!" [4]. The user's silhouette was captured via infrared light and matched to a set of pre-recorded templates. However, the effectiveness of the system as a trainer was not assessed. Becker and Pentland built a Tai Chi trainer [3] in a modified version of the Alive [8] system that provided better depth information. Their system used a hidden Markov Model to interpret the user's gestures and achieved a greater than 90% recognition rate on a small set of 18 gestures. When gestures were recognized but did not receive a high score, the system provided feedback by playing back the segment of the motion in which the student's motion differed most from that of the teacher.

Yang performed some preliminary user studies with his virtual environment called "Just Follow Me" [24] using a *Ghost* metaphor to show the motion of a teacher to a subject. This form of visual feedback is similar to one of our layouts although the rendering styles differed. Subjects executed slow and fast hand motions in one of two test conditions, either while seeing the teacher's motion displayed in a first person view or while viewing the teacher's motion on a third person video display. Although Yang did not quantitatively assess the results, qualitative assessment indicated that the virtual environment produced more accurate trajectories.

3 Training Environment

Our training environment uses the Vicon real-time optical motion capture system from Oxford Metrics. The recorded motion is filtered to reduce noise and recorded for offline analysis. The students wear a wireless HMD through which they see a rendered virtual environment containing animated representations of both the student and the teacher.



Figure 2. The virtual environment for Tai Chi

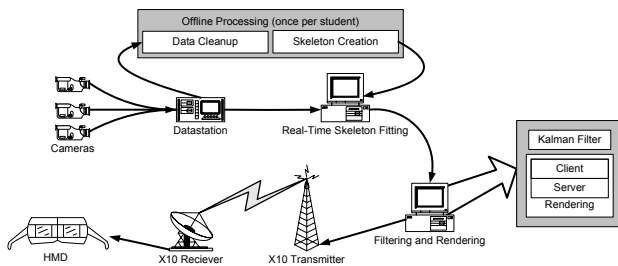


Figure 3. Data pipeline and system diagram

The head-mounted display is extremely light, which is important for a task involving body posture.

3.1 Virtual World

The virtual environment places the student in a pavilion on an island surrounded by a calm lake (Figure 2). On the horizon the sun is setting behind the mountains. The environment is designed to be quiet and peaceful as those are the circumstances in which Tai Chi should be practiced. The pavilion also indicates to the student where the motion capture area is.

In the motion capture laboratory, floor mats are used to demarcate the capture region for the optical motion capture system. To provide further continuity between the real and virtual worlds, the virtual world has a Yin Yang symbol on the floor in the same place that the symbol appears in the real world painted on a floor mat. These physical and visual cues allow the student to easily position him or herself in the appropriate part of the laboratory and virtual environment. This positioning is important because the complete Tai Chi form uses the entire available motion capture area.

The student wears a Spandex unitard with reflective

markers applied to the clothing and skin. The student also wears a fabric hip belt that holds the batteries and the receiver for the HMD. These weigh 3 pounds in total. We use a marker set with 41 14 mm markers, an adaptation of a standard biomechanical marker set. Markers are placed at most joints and along every major bone, and additional markers provide redundant information for critical body segments (head, hands, and feet). When subjects begin the experiment, they are asked to perform a calibration motion in which they move each joint through its full range of motion. The system uses this information to compute a 17 bone skeleton with bones sized to the correct length for that student. The setup process – from having the student put on the unitard through calibration – takes approximately 30 minutes.

3.2 Real-time Optical Motion Capture

Twelve cameras were placed around the laboratory to provide good coverage of a $4\text{ m} \times 5\text{ m}$ capture region 2.3 m high. The lens of each camera is surrounded by infrared emitters that reflect IR light off the retroreflective tape on the markers. The cameras capture 1024×1024 bitmap images at 60 fps. These locations on the image plane of each camera are converted to a set of 3D points by the Oxford Metrics Vicon 512 Datastation hardware.

The locations of the markers are sent to the Oxford Metrics Tarsus real-time server which fits the student's skeleton to the three-dimensional marker locations. The fitting process runs at between 30-60 fps on a 1 GHz Pentium III. The variability in frame rate appears to be caused by occasional spurious marker locations produced by the image processor and large changes in the data.

3.3 Filtering

After the skeleton has been fit to the capture data, the position and orientation of each limb is filtered by the real-time server to reduce noise. Because the fitting process occurs in real time and must be done very quickly, two types of error are apparent in the output. The first is a gross error due to marker mislabeling which causes bones to “flip” out of position. The other error is a general jittering of the skeleton that results from inaccuracy of the fitting process and marker position reconstruction. Flipping errors, which usually occur during high accelerations, are rare in our system since Tai Chi motions tend to be slow. Jittering, however, is noticeable regardless of the speed of the student's motion and needs to be corrected through filtering.

The mean residual error of markers in the fitting process is between 3.9 mm and 7.2 mm for our students (as reported by the Vicon fitting software), enough to cause distracting movement in the virtual environment. In particular when

viewing the environment in first-person, where the virtual camera is fixed to the head's position and orientation, those small errors become large changes in the view.

We chose a Discrete Kalman filter to reduce the noise in our data because it had been successfully applied in systems and applications similar to ours [20]. Sul and his colleagues also used a Kalman filter to post-process motion capture data to smooth the motion and maintain the kinematic constraints of a hierarchical model [16]. We, however, chose to ignore the bone length constraints during the filtering process. This gives acceptable visual results and avoids any extra lag due to the processing time needed to satisfy those constraints.

The Discrete Kalman filter requires the input signal to be in a linear space. This constraint is not a problem for bone position vectors (which are vectors in \mathbf{R}^3), but it is a problem for bone orientations (quaternions representing points in \mathbf{S}^3). In order to apply this filter to the orientation data we first linearize the orientation space using a local linearization technique [7].

Because the virtual camera is fixed to the position and orientation of the student's head, jitter that may have been unnoticeable on other parts of the body is magnified and results in an unacceptably shaky first-person view. To correct this, we apply a more aggressive set of filter parameters to the head position and orientation than to the rest of the body. These parameters make the first-person point of view stable without sacrificing agility in the rest of the body.

3.4 Rendering

We chose the Littech 3.1 game engine for our rendering because it was a fast, robust, and proven engine that supported rendering features critical to our Tai Chi environment and future environments built with our system: smooth skinned models, custom render styles, and efficient networking.

Our rendering application is structured with a client/server model. The server acts as an intermediary, broadcasting commands to control the layout of students and teachers in the environment and the playback position and speed. The server also records data for both the virtual teacher and student for offline analysis. The client takes the data from the filtering application, applies the position and orientation data to an avatar skeleton, and renders the virtual world with the teacher and student models. The client, server, and the filtering process all reside on a 733 MHz Pentium II with a GeForce3 graphics accelerator. The rendering speed averages 17 fps, and is always at least 14 fps.

After the scene is rendered, a scan converter changes the signal from VGA to NTSC. The signal is then sent to an off-the-shelf X10 wireless video transmitter that broadcasts

the video from a location 3 m above the student's head in the center of the capture area. This signal is received by an X10 video receiver on the student's belt, and then sent to the Olympus Eye-Trek glasses we use for our HMD. The rendering is bi-ocular, with the same image sent to each eye via separate LCD screens. Because a lightweight HMD is required for this application, we had to compromise on spatial resolution of the display itself. The Eye-Trek had a 37.5° horizontal and 21.7° vertical field of view, each screen was 640×360 pixels. The weight of the apparatus was 3.4 oz. Therefore, our decision to degrade VGA to NTSC in order to broadcast it implied little or no further loss of resolution. An important detail in rendering the image for the student is to eliminate other light sources in the laboratory so that the virtual environment is the all the subject sees. We accomplish this by turning off the lights in the laboratory rather than via a cowl.

3.5 Latency

The latency of the motion capture and rendering process was 170 ms. To measure our end-to-end system latency, we put markers on a rotating disc revolving at 60 RPM and then added representations of those markers to our existing Tai Chi environment. We then projected the virtual disc image onto the physical disc as it rotated and took photographs to measure the angle between the projected and real markers. The angle difference $\Delta\theta$ between the markers is proportional to latency such that $latency = k\Delta\theta$ where $k = 60s / (60RPM * 2\pi) = 1/2\pi \approx 0.159$. Our measured $\Delta\theta$ was 1.1 radians, giving us an end-to-end latency of 170 ms.

A latency of 170 ms is high for an interactive virtual environment, but the effect on the student's performance is somewhat mitigated by the slow nature of movements in Tai Chi. In addition, the students tend to lag behind the teacher by an average of 610 ms as described in section 4.2.

4 Assessment

In order to evaluate the effectiveness of our system for learning Tai Chi, we tested 40 volunteers to analyze how well they learned the motions of a previously recorded teacher, given different visual user interfaces involving the teacher and student in the virtual space.

There are a large number of design choices for representing the student and teacher in a full body motion training task. We made the constraining assumptions that the student would always see a first-person view, that the teacher and student would be represented in a fairly realistic way (both would remain humanoid, be of the correct proportions, etc.), and that the feedback that the students received would be entirely visual and based on body positions alone. We also fixed the orientation of the teacher and student so

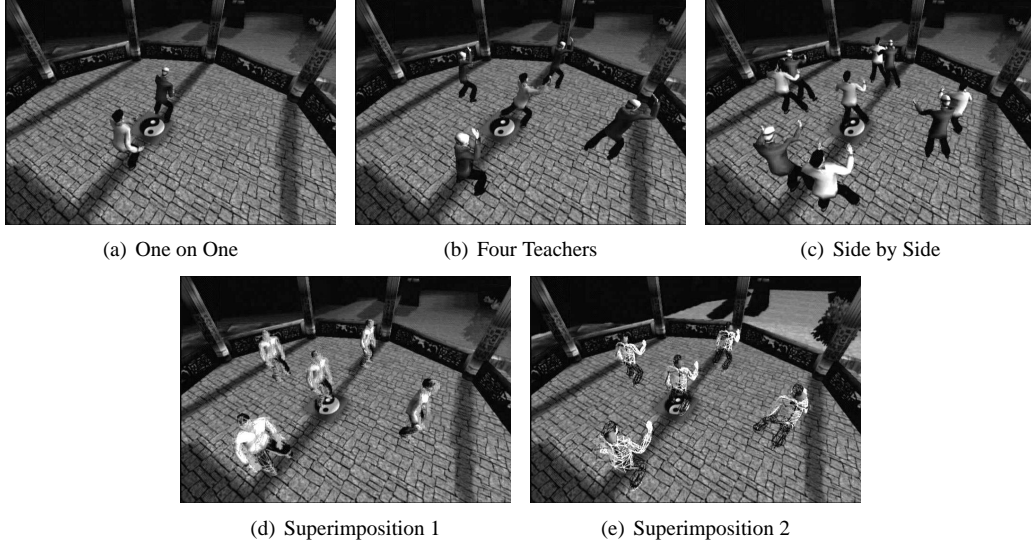


Figure 4. The five conditions

that they faced that same direction. This decision was based both on early pilot studies and the fact that traditional Tai Chi classes are taught this way. After a number of pilot experiments, we identified five representative points in the design space (Figure 4).

- **One on One:** One teacher with one student
- **Four Teachers:** Four teachers surrounding one student
- **Side By Side:** Four teachers standing next to four students with one student in the center.
- **Superimposition 1:** Five normally rendered students with a red wireframe teacher superimposed on them.
- **Superimposition 2:** Five wireframe and transparent green students with a red stick figure teacher superimposed inside them.

4.1 Experiment Design

The majority of our 40 volunteers were college students with no previous formal motion training. Each subject experienced four randomized (layout, motion) pairs where the layout was one of the five in Figure 4 and the motion was one of the four described below. There were originally only four layouts; the fifth, Superimposition 2, was added late in the experiment. Based on early subject feedback regarding the rendering of Superimposition 1, we chose to test an alternate rendering method with the same general layout. No subject repeated a layout or a motion. The order in which the students saw the layouts and motions was randomized to minimize any learning effects between motions.

For each (layout, motion) pair, subjects were asked to match the teacher as they performed twelve repetitions of the same motion after watching the motion from any viewpoint the subject chose. Subjects were told when six repetitions remained and when four repetitions remained.

All four motion animations were segments of the same Tai Chi form performed by a Tai Chi teacher and captured in an offline motion capture session. Each motion was approximately twenty seconds long and had distinguishing characteristics. Motion 1 featured simple hand movements and a 90 degree turn to the right; motion 2 had little hand motion, a 180 degree turn of the feet, and some turning of the upper body; motion 3 featured slow hand movements, large movement of the feet, and some turning of the thorax; and motion 4 featured swift hand movement and turning of the thorax and hips but little movement of the feet. Based on participants' error measurements, we can say that all motions were of approximately equal difficulty except for motion 1 which was significantly easier (Figure 5(b)).

4.2 Results

In order to measure how closely the student mimicked the teacher's motion we used 12 bones for comparison: the upper and lower arms, hands, upper and lower legs, and feet. Each bone was defined using the skeleton's hierarchy with a parent end and a child end. For example the upper arm had a parent end corresponding to the shoulder and a child end corresponding to the elbow. We measured the error for each bone as the difference between teacher and student child end positions, once the parent end positions and bone lengths had been normalized. Given student bone

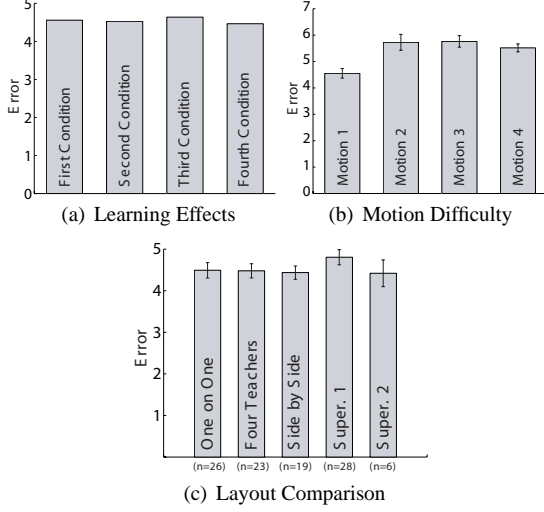


Figure 5. Error analysis results. Error bars show 95% confidence interval.

end positions S_p and S_c and teacher bone end positions T_p and T_c the error E for that bone was:

$$E_{limb} = \left\| \frac{S_c - S_p}{\|S_c - S_p\|} - \frac{T_c - T_p}{\|T_c - T_p\|} \right\|$$

We then average the errors over all limbs to get the total error for one frame of data: $\bar{E}_{frame} = \sum_{i=1}^{12} E_{limb(i)} / 12$. To obtain the total error for one repetition of the motion, we summed over the n frames in that repetition: $\bar{E}_{rep} = \sum_{j=1}^n \bar{E}_{frame(j)} / n$.

Finally, though there are 12 repetitions for each motion, we considered only the last four that the students performed. Before the 9th trial, subjects were reminded that there were four repetitions remaining and encouraged to do their best on the remaining trials. Averaging over these four gives a reasonable indication of how well the student learned the form. So the total error, for one student, given one layout and one motion, is $\bar{E}_{trial} = \sum_{k=9}^{12} \bar{E}_{rep(k)} / 4$

4.2.1 Data Analysis Details

The error measurement for full motions was slightly more complicated. There were two major sources of error which were common for all students who participated in the experiment: time lag and yaw offset. If the student was not oriented in precisely the same way as the teacher, the error in yaw would add an additional error to all bones in the student's body. Even if the student performed the form perfectly, if there was a few degree offset in the orientation the student would have a large error. To minimize this, we considered the student's starting pose and found the yaw offset that minimized this static pose error. We found that




Figure 6. Change in error with yaw and time shifting

when considering all possible offset angles, from -30° to 30° there was one clear point where error was minimized (Figure 6(a) shows a typical case).

In addition to the yaw shift, as the students performed the motion they all tended to lag behind the master. Because they were still doing the motion correctly, just not at the same time as the master, we searched for a time shift that would minimize this error. For each repetition of the motion, we considered all possible time shifts from 0 to 120 frames (0 to 2 seconds). For each value, we compared the shifted student data with the original teacher data to get the error value. Again we found that there was one clear shift value which minimized the error (Figure 6(b) shows a typical case). The average time shift was approximately 610 ms.

We also needed to make sure there was no learning or fatigue effect. Figure 5(a) shows that when we averaged all conditions performed first, second, third, and fourth across all subjects, the order of performance had no effect on error. Before we could average results from different motions we needed to compensate for different difficulties in the motions. Figure 5(b) shows that for the average across all subjects, motion 1 was significantly easier than the other motions. For our final analysis, we normalized the error values based on the mean difficulty of each motion.

Having dealt with these issues, we can now compare all of the layouts with each other. Figure 5(c) shows the average error across all subjects for each layout. The N varies for each of the layouts because Superimposition 2 was added later in the experiment and some subjects were unable to do all four (layout, motion) tests due to lack of time, discomfort, and other reasons. The only statistically significant conclusion we can draw based on the 95% confidence intervals for the average errors is that Superimposition 1 is worse than the others by a small amount. One on One, which most closely mimics a traditional teaching environment, is as good as the other techniques.

4.2.2 Questionnaire Results

After each experiment ended, we gave a questionnaire to the subjects. We asked them to rank the four layouts they expe-




Figure 7. Questionnaire Results: Layout Difficulty

rienced, from easiest to hardest. As Figure 7 shows, the two Superimposition trials were considered to be significantly harder than the other trials. In fact, all of the subjects who tried Superimposition 2 thought it was the most difficult. Interestingly, although subjects considered Superimposition 2 very difficult compared to the other layouts, average error on that layout was not significantly greater than the other non-superimposed layouts.

5 Conclusions

Our primary contribution is the engineering of a wireless virtual reality system that tracks 41 points on the human body in a working volume of 4 meters by 5 meters by 2.3 meters high. Our system has a positional accuracy of roughly 5 mm, a latency of 170 ms, and maintains a rendering frame rate of 17 fps using a commodity PC. Using a light sunglasses-style HMD and a 3 pound waist pack, ours is the first wireless virtual environment system useful for a large variety of applications.

To our surprise, none of the layouts of students and teachers had a substantial effect on learning a motor control task involving full body motion. This result contradicts the initial intuition of virtual reality researchers with which we have discussed this project. They generally assumed superimposition would make the task much easier. If other studies come to similar conclusions for other motor control tasks, the virtual environment research community may need to more carefully examine which tasks are best suited for immersive training, as full body motion motor tasks are often cited as example applications.

We have several theories about why the VR-only tech-

niques fared no better than our simulation of real-world instruction (the One on One layout). In our experiments, students were expected to perform motor movements while simultaneously watching the movements of the teacher's avatar. Other studies of simpler movements have indicated that this simultaneous activity may interfere with learning [12, 14]. Pomplun and Mataric compared the performance of subjects making arm movements when the subjects were instructed to rehearse while watching a video presentation of the motion and when they were instructed to just watch the video before performing the task [12]. The subjects who did not rehearse performed substantially better.

Ellis et al. showed that system latency is highly correlated with performance in a virtual reality tracking task [5] and Ware and Balakrishnan show that mean time for static target acquisition is directly proportional to latency [19]. Although our latency is admittedly large, we chose a slow task so that latency effects would be small. We also measured our errors with a technique that should account for the students' latencies in responding to the teacher. It would have benefited our experiment to understand how our latency – and by extension other factors such as field of view and frame rate – affected performance, and what amount of latency is acceptable for a motion training task such as ours.

With the exception of the superimposition layouts, our rendering styles were chosen to represent the student and teacher in as realistic a fashion as possible given the limited polygon count possible in a virtual environment. However, experimental evidence indicates that subjects tend to track the hands when watching or imitating a teacher performing arm movements [9]. If these results generalize to Tai Chi, a rendering style that emphasized the position and orientation of the hands might be more effective.

We evaluated the motion of the students during trials 9-12 in the virtual environment. A better evaluation technique would have been to remove the visual stimulus of the images of the teacher and student and assess the student's unaided performance [14]. In addition to assessing retention in the absence of feedback, this assessment technique would have allowed us to compare the effectiveness of more standard visual presentations such as a video of a teacher projected on the wall of the laboratory and perhaps understand whether the small field of view and low resolution of the HMD were affecting the student's ability to learn in the environment.

We can not conclusively rule out that our virtual environment was at fault; one can always worry that if only our fidelity had been better, the latency had been lower, or we had just been clever enough to try another rendering style or user interface, the results would be different. However, our system is substantially better in both tracking area and spatial resolution than most full body tracking systems, and

addresses what has been identified as one of the biggest remaining problems: the umbilical cord [18]. Therefore, we can at least conclude that if virtual environments are to be helpful for this task, they must either be at higher quality than we have engineered, or utilize better user interaction techniques than we devised, or both.

References

- [1] Interacting with virtual humans through body actions. *IEEE Computer Graphics and Applications*, 18(1):8–11, Jan.-Feb. 1998.
- [2] N. I. Badler, M. Hollick, and J. Granieri. Real-time control of a virtual human using minimal sensors. *Presence*, 2:82–86, 1993.
- [3] D. Becker and A. Pentland. Using a virtual environment to teach cancer patients T'ai Chi. In *Proceedings of International Conference on Automatic Face and Gesture Recognition*, 1996.
- [4] J. W. Davis and A. F. Bobick. Virtual PAT: A virtual personal aerobics trainer. In *Workshop on Perceptual User Interfaces*, 1998.
- [5] S. R. Ellis, B. D. Adelstein, S. Baumeler, G. J. Jense, and R. H. Jacoby. Sensor spatial distortion, visual latency, and update rate effects on 3D tracking in virtual environments. In *Proceedings, VR '99*, pages 218–221, 1999.
- [6] T. S. Julien and C. Shaw. Firefighter training virtual environment. In *SIGGRAPH Conference Abstracts and Applications*, page 183, 2002.
- [7] J. Lee and S. Y. Shin. General construction of time-domain filters for orientation data. *IEEE Transactions on Visualization and Computer Graphics*, 8(2):119–128, 2002.
- [8] P. Maes, T. Darrell, B. Blumberg, and A. Pentland. The ALIVE system: Wireless, full-body interaction with autonomous agents. *Multimedia Systems*, 5(2):105–112, 1997.
- [9] M. J. Mataric and M. Pomplun. Fixation behavior in observation and imitation of human movement. *Cognitive Brain Research*, 7(2):191–202, 1998.
- [10] T. Molet, R. Boulic, and D. Thalmann. A real-time anatomical converter for human motion capture. In *EGCAS '96: Seventh International Workshop on Computer Animation and Simulation*. Eurographics, 1996.
- [11] J. Molnar, R. Anschuetz, C. Jones, and P. Kelly. Immersion of a live individual combatant into a virtual battlespace. In *The 14th Annual AESS/IEEE Dayton Section Symposium*, pages 27–34, 1997.
- [12] M. Pomplun and M. J. Mataric. Evaluation metrics and results of human arm movement imitation. In *Proceedings, First IEEE-RAS International Conference on Humanoid Robotics (Humanoids-2000)*, September 2000.
- [13] J. Y. Ruisseau, L. Todeschini, and P. Gorzerino. Learning and training simulator virtual reality development for parachute descent. In *Proceedings of the IEA 2000/HFES 2000 Congress*, pages 527–529, 2000.
- [14] R. A. Schmidt and G. Wulf. Continuous concurrent feedback degrades skill learning: Implications for training and simulation. *Human Factors*, 39(4):509–525, 1997.
- [15] S. Semwal, R. Hightower, and S. Stansfield. Mapping algorithms for real-time control of an avatar using eight sensors. *Presence*, 7(1):1–21, 1998.
- [16] C. Sul, S. Jung, , and K. Wohn. Synthesis of human motion using Kalman filter. In *CAPTECH*, pages 100–112, 1998.
- [17] D. L. Tate, L. Sibert, and L. T. King. Virtual environments for shipboard firefighting training. In *Proceedings of the 1997 Virtual Reality Annual International Symposium*, 1997.
- [18] M. Usoh, K. Arthur, M. C. Whitton, R. Bastos, A. Steed, M. Slater, and F. P. Brooks. Walking > walking-in-place > flying, in virtual environments. In *Siggraph 1999, Computer Graphics Proceedings*, pages 359–364, 1999.
- [19] C. Ware and R. Balakrishnan. Reaching for objects in VR displays: lag and frame rate. *ACM Transactions on Computer-Human Interaction*, 1(4):331–356, 1994.
- [20] G. Welch and G. Bishop. An introduction to the Kalman filter. In *ACM SIGGRAPH 2001 Course Notes*, 2001.
- [21] G. Welch, G. Bishop, L. Vicci, S. Brumback, K. Keller, and D. Colucci. High performance wide-area optical tracking: The HiBall tracking system, 2001.
- [22] C. J. Wilson and S. K. Datta. Tai Chi for the prevention of fractures in a nursing home population: An economic analysis. *Journal of Clinical Outcomes Management*, 8(3), 2001.
- [23] C. R. Wren, A. Azarbayejani, T. Darrell, and A. Pentland. Pfinder: Real-time tracking of the human body. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):780–785, 1997.
- [24] U. Yang. Just Follow Me: An immersive VR-based motion training system. In *International Conference on Virtual Systems and Multimedia*, 1999.