

Aaron Bloomfield*

Department of Computer Science
The University of Virginia
151 Engineer's Way
Charlottesville, VA 22904

Norman I. Badler

Department of Computer and
Information Science
University of Pennsylvania
Philadelphia, PA 19104

Virtual Training via Vibrotactile Arrays

Abstract

What is often missing from many virtual worlds and training simulations is a physical sense of the confinement and constraint of the virtual environment. We present a method for providing localized cutaneous vibratory feedback to the user's right arm. We created a sleeve of tactors linked to a real-time human model; the tactors activate to apply sensation to the corresponding body area. The hypothesis is that vibrotactile feedback to body areas provides the wearer sufficient guidance to assume correct body configurations and ascertain the existence and physical realism of access paths. We present the results of human subject experiments that study both explicit and implicit training of skills using vibrotactile arrays. Implicitly, collision awareness is achieved by activating the appropriate tactor when a body part collides with the scene; thus, the user will attempt to correct his or her body configuration. Explicitly, we use the tactors to guide the body into the proper configuration. The results of human subject experiments clearly show that the use of full arm vibrotactile feedback improves performance over purely visual feedback for navigating the virtual environment, as well as allowing easy acquisition of new skills. These results validate the empirical performance of this concept.

I Introduction¹

The virtual experience of a confined space or a training simulation may be satisfying in a visual but not haptic sense. By flying through the space one can get excellent visual impressions of its shape and relationships. Yet there are important reasons for going beyond visual realism. If one needs to reach or maneuver inside the space—say to do complex equipment maintenance or repair—then the flying eye and the disembodied hand are no longer adequate paradigms for virtual equivalents of physical presence in the scene. Thus, applications that require experiencing and testing the feasibility of physical access for the entire body cannot be supported in realistic (i.e., low) cost virtual environment configurations. One such application is maintenance and repair, as learning techniques on a real device exposes it to additional wear and tear, and may require taking it off-line. Another application is the learning of new skills that require fairly precise physical body posture. While qualified instructors can often be found, an initial set of baseline skills can be learned with the use of a

*Correspondence to aaron@virginia.edu.

1. Some of the research discussed in this article was previously presented in Bloomfield and Badler (2007).

simulator with haptic feedback. If the appropriate hardware were readily available, both complex repairs and physical skills could be practiced first in a virtual environment.

Currently, the end-effector haptic experience of the PHANToM device is one of the few practical devices that provide a sense of confined immersion. To feel fully immersed in a virtual environment, haptic feedback should extend beyond the end effector. In the real world, one feels tactile feedback across the entire skin surface, and not just as a force applied to the hands. Simulating this sense of full-body tactile feedback ought to increase the user's sense of presence in the virtual world.

Moving the point of view in a virtual environment is easily accomplished by direct sensing of head position or via interactive input device surrogates such as mice or 3D widgets. Likewise, observing one's body pose through an avatar in the virtual environment is relatively easy with commercially available motion capture systems. While obtaining spatially-limited force feedback at an end-effector is relatively straightforward, for example, with a PHANToM device, supplying a user with force or tactile feedback on a larger body area is very difficult. Typically an exoskeleton approach, which is expensive, clumsy, relatively non-portable, and possibly even dangerous to wear, is the only method (Burdea, 1996).

As an alternative to an exoskeleton, we propose using a tactile actuator (tactor) instead of force feedback. A tactor is a small pneumatic or electronic controlled pressure point that can be actuated as a cutaneous stimulus. The tactor lowers the complexity as well as the cost of the system, while still providing sufficient haptic feedback to allow the user to ascertain sensation on a body part. One highly successful application of tactors is in a flight jacket that is worn next to the skin of a pilot, which provides spatial, gravitational, or threat information during flight (Knapp, 2001; Rupert, 2000). There are many possible tactor designs available today, including voice coil motors, shape memory alloys, or piezoelectric benders (Burdea, 1996; Fletcher, 1996); we used vibratory motors.

In our system, motion capture is used to find the body pose in 3D space; the tactor locations are com-

puted from the body pose and the known positions of the tactors on the user. In the collision awareness modality, intersections between tactor locations and virtual object geometry are checked hierarchically with bounding volumes then with detailed object geometry to optimize performance. Once a collision occurs, the corresponding tactor is energized. The vibrotactile sensation is meant to augment a visual view of the scene and aid the user in establishing a collision-free pose in the space. In the skill acquisition modality, the user is guided to the correct body pose by stimulation of the appropriate tactor(s).

All the experiments described in this article were approved by the Institutional Review Board of the Office of Regulatory Affairs of the University of Pennsylvania (protocol number 708331).

2 Related Work

Only the most relevant related work is presented here; a more complete list can be found in Bloomfield and Badler (2007).

Rupert studied the use of tactor arrays for haptic feedback with the goal to better inform pilots of their aircraft's state. The system used an array of tactors in a flight jacket to signal the true gravity vector (among other aspects) to the pilot's torso during complex flight maneuvers (Knapp, 2001). Similar research has been performed by van Erp (2000), and by Cholewiak with sensory-impaired (blind and deaf) individuals (Cholewiak & Collins, 2000; Cholewiak & Wollowitz, 1992). None of the known systems used tactor arrays to provide body pose collision awareness.

Burdea (1996) provides a good overview of how human nerve sensors cease sending signals to the brain when exposed to a constant stimulation. We analyzed skin habituation in separate experiments to ensure it would not be encountered in our experiments.

Burdea also discusses the two-point limen, which is the minimum distance that various parts of the body can differentiate two distinct points (for example, the points of a draftsman's compass) from a single point. Weinstein (1968) did an extensive experimental study of the two-




Figure 1. *Tactor suit.*

point limen distance, which we used to decide how and where to place the tactors. Although further research has provided alternate means of measuring the limen (Craig & Johnson, 2000), this did not affect our tactor placement.

Researchers have also focused on the efficacy of different types of tactile sensation (Biggs & Srinivasan, 2002). We chose vibratory stimulation because it could be reproduced easily in other haptic environments, it is relatively inexpensive, and the components are easily obtainable.

Bach-y-Rita et al. (1987) define sensory substitution as “the provision to the brain of information that is usually in one sensory domain . . . by means of the receptors . . . of another sensory system . . . examples include sign language for the deaf, and Braille for the blind.” Kaczmarek et al. (1991) present a good overview of the various factors that need to be considered for sensory substitution systems; these were taken into account when designing our system.

Prior research has found that providing haptic feedback will improve a user’s task training performance in a virtual environment. Adams et al. (2001) used force feedback in a VR environment for constructing a LEGO airplane, and found an increase in the user’s performance. Cheng et al. (1996) substituted vibratory feedback for force feedback, but not for collision detection.

Neither of these domains contained vibrotactile feedback for collision detection.

3 Hardware

We constructed vibratory tactors from commercially available parts. We attached a direct current motor with an eccentric mass to a Velcro mount. When a voltage is applied, a vibratory sensation is felt. Our initial use of Velcro straps to attach the tactors to the arm allowed for a very flexible design, but the tactors’ vibrations were transmitted around the arm, and created too diffuse a sensation. Thus, we created three customized tactor sleeves made from long-sleeve, skin-tight elastic athletic shirts of various sizes (Bloomfield, 2003; Bloomfield & Badler, 2003b). The Under Armour shirt (Turf Shirt, item 0032) is thin enough to transmit the vibrations yet strong enough to affix the Velcro mounts. The shirt held the tactors flat against the subject’s skin, allowed freedom of movement, and stretched to fit many different sized people. Figure 1(a) and 1(b) show the suit without and with the tactors and motion capture markers attached.

The tactors were located in four “rings” of four tactors each, which were arranged along the forearm and upper arm, as shown in Figure 2(a). The required colli-




Figure 2. Arm and hand models with tactor placements.

sion awareness task was designed to provide more sensation to the forearm than the upper arm; thus the former had three rings of tactors, and the latter had one ring. Eight tactors were placed on the hand; four around the base of the hand (level with the palm), and on each of the digits except the little finger, as in Figure 2(b). The tactors required by the skill acquisition task were a proper subset of those required for collision awareness.

The hand model in Figure 2(b) was not intended to be realistic; rather it was designed to provide sufficient feedback to the various parts of the hand and lower the computational load on the real-time collision detection routines. The little finger was intentionally removed, as it was not supplied with a tactor.

For the experiment, the subject was placed in a room-sized wireless (infrared) motion capture device (a ReActor by Ascension Technology Corporation). One of the walls of the motion capture device consisted of a large projection screen, which provided the visual display.

The tactors were activated by a series of three relay boards. Each relay board was only capable of switching the power on or off—voltage regulation was not possible. Thus, the amplitude of the tactors' vibration was fixed, but their frequency was modified, as discussed below.

The latency from the visual display was approximately 43 ms, and the latency from the haptics hardware was approximately 68 ms (Bloomfield, 2003). Thus, the visual display updated almost immediately, and the haptic sensations followed shortly (25 ms) thereafter. These

latencies are considered low and medium latency, respectively (Meehan et al., 2003). No subjects reported sensing a delay between the visual alerts and the haptic sensation, based on informal questioning.

4 Collision Awareness Experiments

The original concept when designing this system was to provide a means to substitute tactile feedback for force feedback in a confined space in a virtual environment. Pursuant to this goal, we designed a series of collision awareness experiments to evaluate subject efficacy in sensing the virtual world through the addition of tactile feedback. A number of experiment sets were run; each iteration of experiments allowing further modifications for successive ones. We present here only the most recent set; earlier versions can be found in Bloomfield and Badler (2003a).

4.1 Experimental Design

The experiment consisted of the subjects reaching their right arms into each of the six puzzles shown in Figure 3(a). The objective was to touch the sphere at the end of each puzzle with the virtual hand, without colliding with the rest of the puzzle. Touching the sphere would signal a successful puzzle completion; the puzzle would then disappear, and the subject would




Figure 3. Collision experiment simulation task.




Figure 4. The virtual arm.

move on to the next one. The puzzles could be completed in any order. If a collision did occur (other than the hand colliding with the sphere), the puzzle would not complete. This ensured that the subjects made an effort to minimize the number of collisions.

A partially completed experiment is shown in Figure 3(b). Although the puzzles look small in the figure, they were displayed on a 2 m high by 2.5 m wide (6' \times 7.5') projection screen, which the subject was facing. To reduce the chance of simulator sickness, the perspective was fixed, and did not change as the subject moved (Kolasinski, 1995). Thus, as the subject moved his or her arm, a disembodied arm and shoulder was displayed moving about the virtual environment. The arm was

shown translucent ($\alpha = .5$) so that the subject could see the parts of the puzzle it occluded. Although the translucent arm is difficult to see in Figure 3(b), it is more visible in Figure 4(b). This particular puzzle perspective was chosen so that, in some of the puzzles, it was difficult or impossible to see all of the hand when it was inside the puzzle.

The puzzles required the arm be inserted up to about the elbow in order to touch the sphere, as shown in Figure 4(a). Note that the puzzles are shown translucent and the arm opaque to create this image; in the simulation, the reverse was true.

The length of the subject's forearm and upper arm were normalized to a fixed length within the graphical

coordinate system, regardless of the length of the subject's actual arm. This ensured that each subject had to perform the same amount of reach to get into the puzzles.

4.1.1 Feedback Modalities. This experiment provided the subjects two types of feedback: visual and haptic (tactors). Each of the two types of feedback could be either on or off for a particular subject. This created four experimental groups to which the subjects were randomly assigned. The groups are referred to as tactile, visual, none, and both in this article, indicating which of the two feedback modalities they received during the collision experiments.

The arm was shown translucent so the subjects could see the occluded sides. The parts of the arm that collided with an object turned *red*, as shown in Figure 4(b); this was the visual feedback.² The parts of the arm that are not colliding are *blue*. Note that the collisions are colored *magenta* in the figure to increase contrast. The colliding areas (the elbow, the middle of the forearm, and the hand) are a lighter shade of gray in the grayscale version of Figure 4(b).

Note that all groups saw their arm moving in the virtual environment; the visual feedback consisted of seeing the areas of the arm turn *red* that were colliding with a virtual object.

A “deep” collision was indicated by a pulsing tactor (as opposed to a tactor continuously on for non-deep collisions). Subjective and objective results showed little difference, in perception or performance, with the use of the pulsing tactors. Thus, we do not report on them here; for full details, see Bloomfield and Badler (2007).

4.2 Experimental Procedure

Right-handed subjects were recruited from the student and staff population at the University of Pennsylvania. All subjects were fluent English speakers, and all had at least a high school (or equivalent) level of education. Subjects were paid for their participation. The experiment was performed in front of a flat projection

screen. The subjects were shown a video demonstration, with narration, of the task that they needed to perform in the experiment. The narration was intentionally designed to be easy to understand (Flesch-Kincaid reading level³ of 6.1). The subjects were allowed to ask questions about the demonstration. Two questionnaires were given, one after the demonstration and before the experiment, and the other after the experiment.

The subjects were told that all the puzzles pointed straight back, and from this particular perspective the outside puzzles looked to be pointing inward. Confusion over this perspective in prior experiments caused a number of the subjects to angle their arms inward, preventing them from completing the puzzles with the minimum number of tactor activations. The subjects were told that the height of the puzzles was adjustable, and were instructed how to lower them to their preferred height. Lastly, in an effort to reduce the number of objectives for the experiment, the subjects were told that it did not matter how long they took to perform the experiment; the objective was solely to lower the number of collisions.

Subjects were randomly assigned to one of the four experimental groups described above. Subjects who were in one of the two groups that used tactile feedback were given a demonstration of what a constant tactor activation and a pulsing tactor activation felt like, so that they could compare the two. They were told that the pulsing tactor activation was used to indicate deep collisions.

All subjects wore all the tactors, even if the tactors were not going to be activated for their experimental runs. This was done to ensure that arm fatigue would not be a differentiating factor between the experimental groups. For subjects in the two groups that received tactile feedback, the tactors were checked to ensure they were all working after suiting the subject up and prior to running through the simulation.

3. The Flesch-Kincaid reading level is computed by $.39 \times \text{AWS} + 11.8 \times \text{ASW} - 15.59$, where AWS is the average words per sentence, and ASW is the average number of syllables per word. It corresponds to the US grade-school reading level of the given document.

2. Color versions of these figures are available as supplemental files accompanying the online version of this article.




Figure 5. Average factor activations per trial.

4.3 Statistics and Experimental Measures

The main data point was the number of factor activations (described below) that occurred. Each group performed six trials of the simulation, where each trial consisted of completing the six puzzles shown in Figure 3(a). The first two trials counted as training runs, and thus the last four trials were averaged to produce this metric. The results showed that the subjects' performance stabilized after the first two trials, as shown in Figure 5. Time taken to complete the simulation was also measured, but the subjects were told that time taken did not matter. Other metrics recorded included the number of collisions on the arm and hand, and the number of deep collisions.

The number of factor activations per trial is a weighted metric. The total number of factor activations for each frame of the motion capture system (which ran at 33 frames per second) was summed to create this metric. Thus, a single factor active for 1 s would count as 33 activations. For the groups without tactile feedback, the activations were computed as if they had activating factors.

4.4 Experimental Results

A total of 42 valid experimental trials were performed, 25 males and 17 females.

Table 1. Subject Demographics

Group	<i>n</i>	♂	♀	% ♂	Avg age
All	42	25	17	59.5	24.6
None	11	8	3	72.7	25.3
Visual	10	4	6	40.0	24.5
Tactile	11	8	3	72.7	26.5
Both	10	5	5	50.0	22.0

4.4.1 Demographics. The demographics of the valid subjects are shown in Table 1. Demographic data (sex, age, and education completed) were analyzed across all the groups using *t*-tests, and no statistically significant differences were encountered ($\alpha = .05$).

4.4.2 Objective Results. The main metric for the objective results is the number of factor activations, given in Table 2. Only the average of the last four trials was used unless otherwise indicated.

As can be seen in Table 2, the groups with tactile feedback performed significantly better than the groups without tactile feedback. One result that we were not expecting was that the group with only tactile feedback performed better than the group with both tactile and visual feedback. This is discussed below.

The average time taken to complete the trials is also given in Table 2. The subjects were told to take as much time as they needed, and thus we do not use the time taken as a measure of performance. However, it is interesting to compare the times, which is done in more detail below. Statistical *t*-tests were used to analyze the trial runs. Significant differences existed in the total number of factor activations, but not in the average time taken per trial. This is also discussed in more detail below.

A graph of the average number of factor activations per group per trial appears in Figure 5. The average number of activations dropped off significantly after the first two trials, and generally plateaued by the third trial. None of the plotted lines monotonically decrease.

To analyze the length of the learning curve, a one-way ANOVA was run on the number of collisions,

Table 2. Experiment Subject Performance: Tactor Activations and Time

Group	<i>n</i>	Avg tactor activations	Activations <i>SD</i>	Completion time (s)
All	42	3,365.6	2,549.4	91.2
None	11	4,886.1	3,423.4	86.2
Visual	10	4,355.6	1,674.4	92.5
Tactile	11	1,915.8	1,391.9	88.4
Both	10	2,298.1	1,580.3	98.5

where each trial/group combination was an ANOVA group. A Tukey post hoc analysis showed that all the groups showed a significant improvement between the initial trial and either the third or fourth trial (the group with both feedback modalities was significant at $\alpha = .10$, all others at $\alpha = .05$). However, there was no significant change between trials three and six for any of the groups (at $\alpha = .10$). Thus, we proceeded with our initial assumption that the first two trials constituted the learning curve, and the performance generally plateaued for the last four trials. A visual inspection of the graph in Figure 5 further confirms this hypothesis.

An interesting and unexpected feature can be seen in the plotted data of the group without any feedback between trials three and four: the number of activations increased. While all the groups had an increase at some point, the increase for this group is more pronounced. Our hypothesis as to the cause of the decrease in performance is that this group had more arm fatigue from the first two trials than the other groups, as they took longer for the first trial. The decrease in performance (increase of tactor activations) from trials three to four for the group with no feedback was from 3,268.7 to 6,552.1. Because of the wide *SD* for the group with no feedback during trial four, this decrease in performance is not statistically significant. However, the decrease in performance from trial three to trial five (3,268.7 to 5,949.7) is statistically significant (at $\alpha = .05$).

The time taken by each subject was recorded, although they were told that they could take as long as needed. The group with no feedback took longer to perform the first trial than the other groups. This is not

a statistically significant increase in time over the average (at $\alpha = .05$). This supports (but does not prove) the hypothesis that arm fatigue led to decreased performance in trial four.

Other metrics that were analyzed (hand activations, arm activations, deep activations, etc.) did not produce any further interesting results.

4.4.3 Subjective Results. The questions asked on the questionnaires are summarized in Table 3, along with the results. The first five questions were on the pre-experiment questionnaire. Questions 7 to 14 dealt with the tactors, and thus only two of the groups answered them. Likewise, question 6 dealt with the visual collision alerts, and only two of the groups answered that question. Each question was rated on a Likert scale of 1 to 5, where 1 meant “very low” or “very little” and 5 meant “very high” or “very much.”

The amount of 3D experience was the main subjective metric used for comparing the groups, and this is discussed in more detail in the analysis. The exact question was, “How much experience have you had with 3D environments (games, animation programs, graphical programming, etc.)?” There was no statistically significant difference between the 3D experiences of any of the group combinations. The amount of virtual reality experience is a less useful metric—it was much less common to have virtual reality experience than 3D experience, and thus a larger number of subjects would be required before we could effectively compare this result.

All of the questions and their answers are analyzed in

Table 3. Experiment Questionnaire Questions and Results

Question	Feedback type				
	All <i>n</i> = 42	None <i>n</i> = 11	Visual <i>n</i> = 10	Tactile <i>n</i> = 11	Both <i>n</i> = 10
1. Amount of prior 3D experience	2.76	2.45	2.50	3.18	2.90
2. Amount of prior VR experience	1.74	1.36	1.70	1.91	2.00
3. Motion sickness susceptibility	1.45	1.36	1.80	1.36	1.30
4. Computer eye strain susceptibility	1.36	1.36	1.30	1.45	1.30
5. Perceived difficulty of demo actions	2.07	2.18	2.30	1.82	2.00
6. Collision realization aided by visual feedback	3.40		3.60		3.20
7. Collision realization aided by tactile feedback	4.57			4.27	4.90
8. Comfort of wearing the tactors	3.43			3.55	3.30
9. Amount that skin saturation was encountered	1.52			1.45	1.60
10. Tactor was a signal (1) or spatial indicator (5)	2.67			2.64	2.70
11. Accuracy of the tactor activations	4.24			4.36	4.10
12. Intensity of a pulsing tactor vs. a constant tactor	3.43			3.73	3.10
13. Ability to sense pulsing tactors	2.52			2.36	2.70
14. Amount the pulsing tactors helped	2.48			2.27	2.70

Bloomfield and Badler (2007); we only report the more interesting results here.

The tactile feedback seemed to help subjects more than the visual feedback (average responses of 4.57 versus 3.40). This agrees with the objective results, described below, which show that the groups with tactile feedback performed better than the groups with no tactile feedback.

Subjects felt they encountered very little skin sensation saturation (average response of 1.52), as was indicated by previous skin habituation experiments.

Question 11 had some disappointing results. The exact wording of the question is, “How much did you feel the tactor acted only as a signal (telling you that there was a collision somewhere, but not where that collision was) as opposed to a spatial indicator (telling you exactly where the collision occurred)?” The allowed responses ranged from only as a signal for 1, to both as a signal and a spatial indicator for 3, to only as a spatial indicator for 5. The result, 2.67, was less than we had expected. It was consistent across the two groups that

used the tactors. However, the objective results described below show a significant improvement with the use of the tactors. Thus, while the subject’s intuition may have caused him or her to rate the tactor activations more as a signal, it may still have served quite well as a spatial indicator. The fact that multiple tactors were activated at any given time also contributed to the lower than expected result for this question.

4.5 Analysis

A one-way ANOVA was performed on the average number of activations for the last four trials for each group ($F(3, 12) = 6.989, p = .0057$). A Tukey analysis indicated that the greatest differences existed between the tactile group and the group with no feedback, and between the group with both feedbacks and the group with no feedback. A one-way ANOVA for time did not show any statistically significant differences between the groups ($F(3, 12) = 0.6085, p = .62$).

Statistical *t*-tests were run on the average number of

Table 4. *t-Test Tactor Activation Analysis Results*

Group feedback	Group feedback	ν (DOF)	t value	Significance (α)
None	Visual	19	0.457	.321
None	Tactile	20	2.666	.007
None	Both	19	2.257	.014
Visual	Both	18	2.826	.006
Tactile	Both	19	-0.586	.275

activations for the last four trials for each group. The results of the t -tests are shown in Table 4. The last column shows the exact significance level (α value) for that row. This table contains the main data analysis that supports our hypothesis. Note that it does not make sense to compare the remaining group combination (visual feedback versus tactile feedback) via t -tests, as one cannot be considered a control for the other.

In the three combinations where the difference between the two groups included the addition of tactile feedback (the middle three rows), the results were all significant (at $\alpha = .05$). Furthermore, in the two combinations where the only addition was the tactile feedback (the second and fourth rows), the results were highly significant (at $\alpha = .01$). For the two groups where the difference between the two groups did not include the addition of tactile feedback, and thus was only the addition of visual feedback (the first and fifth rows), the results were not significant (at $\alpha = .05$).

There is also a significant increase from the group with no feedback to the group with both feedbacks (the middle row). While it is significant at $\alpha = .025$, it is not as significant an increase in performance as with the two combinations where the only addition was the tactile feedback (the second and fourth rows).

We were expecting the increase in performance that we found with the addition of the tactile feedback. What we were not expecting was the decrease in performance from the tactile only group to the group with both feedbacks. This difference in performance, be-

tween the tactile only group and the group with both feedbacks (the bottom row in Table 4) is not significant. Our hypothesis is that the addition of both feedback modalities added too much information, causing the subjects to be distracted by the visual alerts, and to not respond to the tactile alerts as well as the tactile only group. This was supported by a few informal subject comments after the experiments.


There were two questions that had statistically significant differences on the questionnaires (at $\alpha = .05$). The group with no feedback differed significantly with both the tactile feedback group and the group with both feedbacks on question 2 (amount of prior VR experience). The tactile group and the group with both feedbacks differed significantly on question 7 (collision realization aided by tactile feedback). The results from these questions are shown in Table 5, along with the performance results reported above. The questionnaire results that were statistically significantly different are shown in bold in each column.

The differences in the amount of virtual reality experience (question 2) did not concern us. As mentioned above, it is much less common to have VR experience than 3D experience, and thus a larger population would be required before we could effectively compare this result. Because there were no statistically significant differences in the amount of 3D experience, and due to the population size ($n = 42$), we did not feel that the virtual reality experience metric exposed a significant difference between the experimental groups.

The difference in how much the tactors helped be-

Table 5. *t*-Test Statistically Significant Questionnaire Differences

Group	<i>n</i>	Tactor activations	Q2 avg	Q2 <i>SD</i>	Q8 avg	Q8 <i>SD</i>
All	42	3,365.6	1.74	0.66	4.57	0.73
None	11	4,886.1	1.36	0.64		
Visual	10	4,355.6	1.70	0.64		
Tactile	11	1,915.8	1.91	0.67	4.27	0.86
Both	10	2,298.1	2.00	0.45	4.90	0.30

**Figure 6.** Ryukyu Kempo (karate) moves.

tween the groups is very high for both groups (4.27 and 4.90). It is interesting to note that of the two groups with tactile feedback, the group that performed worse (the group with both feedbacks) felt that the tactile feedback was more important than the group with only tactile feedback, perhaps because they judged it based on the visual feedback, and the visual feedback may have hindered them.

5 Skill Acquisition Experiments

One of the many applications of full-body tactile feedback is to teach new physical skills. The previous set of experiments performed training implicitly—the subjects, at the end of the experiment, were trained to perform that task better through feedback when they made a mistake (i.e., collided with the scene). In the second

set of experiments, we studied a more explicit means of training, where the subject's body configuration was oriented to the proper position by indicating which part of his or her right arm needed to move, and in which direction.

The tactor suit used in these experiments provided only right arm feedback, so a viable skill needed to be taught that used only the position and orientation of the right arm. Prior research has discussed teaching martial arts moves through the use of a virtual environment (Kirner et al., 2001), and found that the subject was better able to learn the moves through the use of the virtual reality simulation.

Moves from Ryukyu Kempo, an Okinawan form of karate, are ideal for this experiment, as the four moves selected required only the positioning of the right arm. One of the aspects of Ryukyu Kempo is that the techniques are done with a specific angle of the wrist

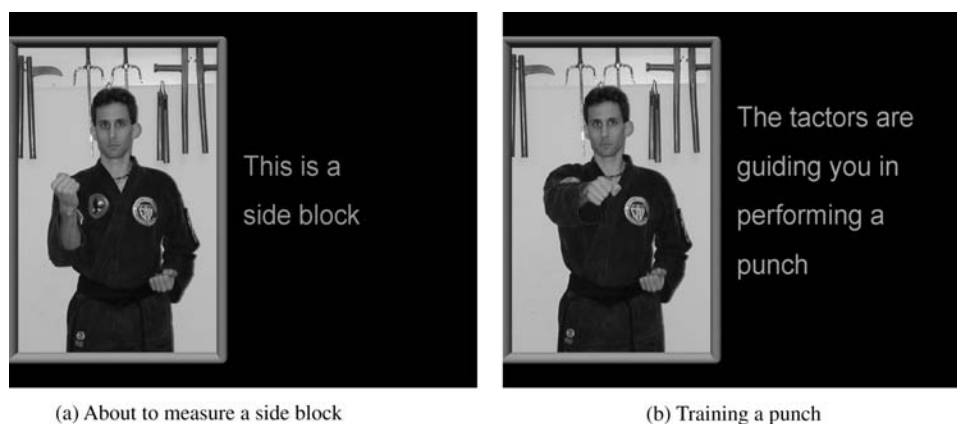


Figure 7. Skill acquisition simulation screen shots.

(Dillman & Thomas, 1992; Terry, 2006). In Figure 6(d), note that the punch is rotated 45° from the horizontal, so that the knuckle of the index finger is on top. All of the moves used in these experiments have this rotation of the wrist. The experiment was designed to teach the correct position and rotation of the arm so as to replicate karate moves. Specifically, the wrist rotation, combined with the positions of the elbow and wrist, were used as the metric to measure subject performance for this experiment.

The first author is a second degree black belt in Ryukyu Kempo, and was the exemplar for the moves described here. Accurate move positions were recorded with the motion capture system, and these arm orientations were then used in the experiment as the “accurate” arm orientations for each move.

5.1 Experimental Design

The subjects were taught, and then required to perform, a series of four martial arts moves, which are shown in Figure 6. Note that Figure 6 is shown with a white gi (karate outfit) to improve contrast on the gray-scale versions of this article.

These experiments were conducted concurrently with the collision experiments described above. The subjects performed these experiments immediately after the collision experiments.

5.2 Experimental Procedure

All subjects were asked if they had any prior martial arts experience. If they did, then they were excluded, as they would either already know the move, or be able to learn it much faster, due to their prior experience with similar martial arts moves.

The experiment itself consisted of a training simulation that lasted about 10 min. The subjects did not see their arm on the screen, but instead saw images and instructions about what to perform for each part of the experiment; this is shown in Figure 7. The simulation went through the training of the four moves in sequence, alternating between directing the subject to perform a move, and guiding him or her into the move’s correct position with the factors.

Each move consisted of five parts. On the first, third, and fifth parts, the subject was asked to perform the move so that a measurement could be taken. For each move, a picture of the move was shown on the screen for the subjects to mimic. One of these pictures can be seen in Figure 7. The subject was asked to hold the position for 15 s, which allowed the simulation to measure the accuracy of the subject’s technique. The method of determining how the accuracy was measured is described below (Section 5.3).

On the second and fourth part, the “factor training sessions,” the factors were used to guide the subject

into the correct position. The shoulder was assumed to be in the correct position, and the rest of the arm's position was determined relative to the shoulder's position. The subject was assumed to be facing the screen, as there is no viable way with the current tactor suit to provide tactile feedback to indicate that the subject needs to be facing a different direction. The most proximal body part was corrected first. Thus, the elbow was first moved to its correct position, followed by the wrist position, followed by the wrist rotation.

The simulation proceeded through each of the five parts for a given move before proceeding to the next move. This prevented the subject from forgetting the trained position between testing, as the testing for a given move was not interspersed with other moves. The order of the moves is the same as the order shown in Figure 6.

Prior to the experiment, the subjects were told about the various sensations they were about to receive, and which sensation meant moving the wrist and which sensation meant rotating the wrist. The sensations were designed to be fairly natural—a single tactor activating on the wrist or elbow meant to move away from that direction; the four tactors on the wrist activating in a rotating manner meant to turn the wrist in the direction of the rotation. The subjects then proceeded to perform the experiment. A screen shot of the simulation is shown in Figure 7. The other screen shots were similar—a move was displayed on the left side, and instructions (such as “please perform a punch”) were displayed on the right side.

The first part for each move provided a base measurement of how accurate the subject's move was. The third part provided a measurement of how accurate it was after one training session, and the fifth provided a measurement of how accurate it was after two training sessions.

5.3 Statistics and Experimental Measures

There were two metrics that were used to measure subject performance: accuracy of the move (distance

from the correct location) and time taken to train the move.

Recognizing martial arts moves is a difficult task, due to the complexity of interpreting human movement and posture (Shinagawa et al., 1997; Sun et al., 2002). For this experiment, we determined the accuracy of the move based on three factors: the positional difference of the elbow and wrist with the correct positions, as well as the scaled angular difference of the wrist rotation with the correct rotation. All locations were relative to the position of the subject's right shoulder. As mentioned above, the subject was assumed to be facing the screen. The subject had to hold each position for 15 s; at 33 motion capture frames per second, the system recorded 495 frames. The simulation took the 33 consecutive frames that had the lowest average distance in that 15 s segment. Thus, this measurement was the 1 s period of time that the subject had the most accurate move.

For purposes of comparison, the forearm and upper arm were each normalized to 3.0 units long from the center of the shoulder to the center of the elbow, and from the center of the elbow to the center of the wrist. The scaling for the angle is $\frac{1}{45}$, so a 45.0° rotation off from the correct rotation would count as a difference of 1.0. Thus, if the subject's elbow was 0.5 units away from the correct position, the wrist was 0.75 units away from the correct position, and the wrist was rotated 15.0° from the correct rotation, the distance would be

$$0.5 + 0.75 + \frac{15}{45} = 1.58.$$

The second metric used was time. For the second and fourth parts for each move, the tactors guided the subject into the correct position and orientation. The tactors were activated proximal (elbow) to distal (wrist), followed by the wrist rotation. If a more proximal joint moved out of position, the simulation would go back to correct that more proximal joint. The time it took until the subject had the correct position and orientation of the arm was measured, and forms this metric.

The simulation did not demand absolute accuracy from subjects, but instead offered positive results as long as the subject stayed within a given range. Specifi-

Table 6. Skill Acquisition Distance *t*-Test Results

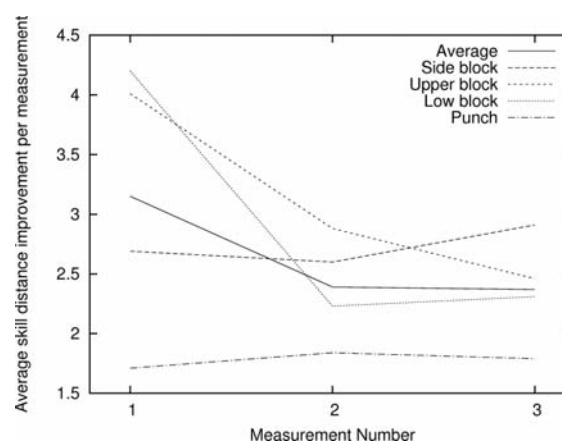
Move	Distance measurements			<i>t</i> -values			Significance at $\alpha = .05$		
	1	2	3	1→2	2→3	1→3	1→2	2→3	1→3
All moves	3.15	2.39	2.37	2.858	0.063	2.565	Yes	No	Yes
Side block	2.69	2.60	2.91	0.170	-0.573	-0.376	No	No	No
Upper block	4.01	2.88	2.46	1.710	0.524	2.143	No	No	Yes
Low block	4.20	2.23	2.31	2.617	-0.124	3.035	Yes	No	Yes
Punch	1.71	1.84	1.79	-0.259	0.090	-0.201	No	No	No

cally, the simulation allowed for a range that the subjects needed to be in for the trainings: the positions needed to be within a distance of 1.0 unit from the correct location. This was determined by a series of test trials. Anything below this value (i.e., more accurate) made it too difficult for the subjects to complete the trial successfully. Once the elbow, for example, is within a distance of 1.0 from the correct location, the elbow tactors will not activate again unless the elbow moves more than 1.25 units away from the correct position. Originally, when the trials were run without this buffer zone, subjects whose positions were right on the border of distance 1.0 from the correct location experienced rapidly alternating wrist and elbow tactors, as when the elbow moved outside the valid distance the wrist correction shut off and an elbow tactor activated. Note that only one tactor was activated at any given time. With the addition of the buffer zone, subjects were better able to retain the elbow position and proceed through the wrist movements.

The position of the wrist is not normalized based on the position of the elbow. Thus, regardless of where the elbow is, the wrist still has only one position that is the correct position at which it is supposed to be (all positions are based relative to the position of the right shoulder). This is true both for the tactor training and the measurements.

5.4 Experimental Results

5.4.1 Demographics. Five subjects participated, three males and two females. The subjects all partici-

**Figure 8.** Skill acquisition distance improvement per trial graph.

pated in this experiment immediately following their participation in the collision experiments.

5.4.2 Objective Results. The objective results are from two metrics: the deviation distance from the correct position, and the amount of time taken for the tactor training.

The distance measures, averaged for all five subjects, are shown in Table 6. As all the *t*-test combinations had the same number of subjects in each group (i.e., 5), the number of degrees of freedom is fixed at $\nu = n_1 + n_2 - 2 = 8$, and the critical value is fixed at $t_{\nu=8, \alpha=0.05} = 1.860$. For each move, there were three *t*-tests that could be performed on the distance: from trial 1 to trial 2, from trial 2 to trial 3, and the overall improvement from trial 1 to trial 3. Each of these is indicated in sepa-

Table 7. Skill Acquisition Time *t*-test Results

Move	Time (s)		Change (s)	<i>t</i> -value	Significance at $\alpha = 0.05$
	1	2			
All moves	18.29	14.81	-3.49	0.482	
Side block	22.79	14.88	-7.91	0.644	No
Upper block	16.62	4.16	-12.45	1.977	Yes
Low block	29.16	25.94	-3.22	0.191	No
Punch	4.60	14.24	9.64	-1.328	No

rate columns in Table 6. A graph of the distance measurements is shown in Figure 8.


The largest improvement can be seen in the upper block and the lower block, the second and third moves performed. With these two skills, the simulation clearly showed an increase in the subject's ability to perform the move with the help of the factor feedback.

The time measurements, averaged for all five subjects, are shown in Table 7. As with the distance measures, all the *t*-test combinations had the same number of subjects in each group (i.e., 5), and thus the number of degrees of freedom is fixed at $\nu = n_1 + n_2 - 2 = 8$, and the critical value is fixed at $t_{\nu=8, \alpha=.05} = 1.860$. The "change" column is the difference in seconds between the first factor training and the second.

A graph of the time measurements is shown in Figure 9.

As can be seen in the graph, all the training times except the punch decreased from the first factor training to the second. Our hypothesis with the increase in training time with the punch is discussed below, in the analysis. Even though the increase in times for the punch from the first factor training to the second is large, the large standard deviation prevented this from being statistically significant. Only the upper block had a statistically significant change (decrease) in the training times. Even with the increase in training times with the punch, the overall training times averaged across all the moves decreased from the first factor training to the second.

5.4.3 Subjective Results. The questions asked of the five subjects who participated in the skill acquisition

**Figure 9.** Skill acquisition training time per trial graph.

experiments, along with their results, are shown in Table 8.

The subjects felt that their elbows were guided into the proper position well (average response of 4.00), but less so about the guiding of the wrist into the proper position (average response of 3.20). Our hypothesis is that the necessity of keeping the elbow in the proper position while moving the wrist into the proper position increased the difficulty and thus lowered the response on the latter.

A number of subjects complained that the wrist rotation signals were too fast or too slow. From the results, we found that the speed chosen (1.2 rotations around the wrist per second) seemed to be about right for most of the subjects. Very little skin saturation was encountered (average response of 1.40), which closely matches

Table 8. Skill Acquisition Questionnaire Results

Num	Question	Average <i>n</i> = 5
17	How well tactors guided elbow and wrist to the correct position	4.00
18	How well tactors guided wrist to the correct rotation	3.20
19	Were the wrist rotation signals too slow (1) or too fast (5)	3.20
20	How well the moves were learned compared to a picture only	4.20
21	Amount that skin saturation was encountered	1.40
22	How well the move guided into matched the move presented	3.80

the results for this question from the collision experiments (1.52).

The subjects felt that the tactor feedback made it much easier to learn the moves than had they only been able to look at a picture (average response of 4.20). Lastly, the subjects felt that the moves that the tactors guided them into matched the presented moves fairly well (average response of 3.80).

5.5 Analysis

The overall improvement for all the moves was statistically significant between the first two measurements, and between the first and third measurements. This, along with the average distance improvements for all the moves (the top row in Table 6), shows that the most learning was achieved during the first tactor training session. The significance of these results, as well as the fact that the averages of the distances showed a significant improvement, supports our hypothesis that skill acquisition can be achieved with the use of vibratory tactors.

The side block and punch did not show significant improvements; indeed, the punch showed an increase in the training time. Our hypothesis on the lack of improvement on these two moves is twofold. One is that both the side block and the punch are fairly natural positions. This is particularly true of the punch, as the hand is essentially held straight out in front of the body. Evidence to support this can be seen by the fact that the distance measurements for the punch are all lower than

for the other moves. Since they therefore came more easily to the subjects on the first trial, there is a limit to the amount of improvement that can be obtained. The second possibility for the side block is that since it was the first move performed, the subjects were still becoming accustomed to the way the tactors directed the arm. Again, this limited the possibilities for improvement.

6 Discussion

In the collision experiments, the result we found most surprising was that the tactile group performed better than the group with both feedbacks (1,915.8 and 2,298.1 activations per trial, respectively). There are a number of possible explanations for the lack of increase in performance for the group with only visual feedback. One explanation is that the virtual environment may have handicapped the visual only group more. The perspective of the puzzles, as well as the transparency of the arm itself, was purposely designed to occlude the subject's view of the inside of the puzzle. However, the highly statistically significant results (the two combinations where tactile feedback was added both had results that were statistically significant at $\alpha = 0.007$) implies that this is not the sole cause of the lack of performance from the visual only group, and that the haptic feedback modality provided a major influence in the performance differences.

An alternative explanation is that the simultaneous collision caused an overload of tactile feedback. During

a collision, all the factors that were colliding with the offending object were activated at once. Although the subjects were obviously able to discern the exit vector direction, as evidenced by their performance results, this method of tactile feedback can be improved upon. Only activating one factor, which would indicate the shortest exit vector, would reduce the amount of tactile saturation and sensation overload the subjects experienced. The fact that this system worked well with this sensation overload indicates that even better results could be obtained if the sensation overload were reduced or removed.

The skill acquisition experiments were meant to be a proof-of-concept experiment, but the positive results that we achieved convinced us to report on them here. While the experimental population was small, we were still able to obtain a significant improvement when using the factors. Further experiments will be needed to further refine the simulation, and to have a direct comparison with people who attempt to learn the skill with the visual aid but without the factors. Future research can examine the creation of more experimental groups with martial arts instructors, which could provide a direct comparison between teachers and factors.

Despite the research on multimodal input modalities, a question remains on how much weight is assigned to each of the modalities in a specific input situation. The results of our research and other studies yield the hypothesis that each type of input modality is best for perceiving a different type of information. The input modality that is best will weight most heavily. In this context, best is a subconscious decision. There is currently no qualitative way to determine a priori what the best input modality is for a given situation. This promises rich avenues for future research.

7 Conclusions

Full body haptic feedback is an eventual necessity for a fully immersive experience in a virtual world. One cannot feel completely “in” a virtual world when only the hand is receiving haptic feedback. In the Star Trek

holodeck, for example, participants feel the physicality of the virtual environment. While we are a long way from applying the physics of actual forces from the external world, cutaneous stimulation can contribute part of this tactile experience. There are many virtual situations that may benefit from haptic feedback; awareness of collisions being only one of them. Refinements on our experiments suggest promising options for future research in full body haptic feedback for improved collision perception and skill acquisition.

The question that motivated this research was whether (whole body) tactile feedback can train an individual in reach and access maneuvers in a virtual confined environment, whether through explicit means or implicit means. Our system used a number of small vibratory factors on the subject’s right arm and hand. Coupled with controlling hardware and simulation software, the result is a fully immersive simulation where the subjects feel collisions with virtual objects through vibrations applied to their skin. Through a large set of formal human subject experiments, the resulting data for the collision experiments clearly show a significant reduction in virtual collisions in the subjects who used the factors over those who did not. Further skill acquisition experiments showed a positive increase in body configuration through use of the factors.

Acknowledgments

This research was funded in part by NASA and the Office of Naval Research.

References

- Adams, R. J., Klowden, D., & Hannaford, B. (2001). Virtual training for a manual assembly task. *Haptics-e*. Available at <http://www.haptics-e.org> 2(2).
- Bach-y-Rita, P., Webster, J. G., Tompkins, W. J., & Crabb, T. (1987). Sensory substitution for space gloves and for space robots. In *Proceedings of Workshop on Space Telerobotics*, 2, 51–57.
- Biggs, J., & Srinivasan, M. A. (2002). Tangential versus nor-

- mal displacements of skin: Relative effectiveness for producing tactile sensations. *Proceedings of the 10th International Symposium on Haptic Interfaces for Virtual Environments and Teleoperator Systems*, 121–128.
- Bloomfield, A. (2003). *TRACE: Tactor reach access and constraint environment*. Doctoral dissertation, University of Pennsylvania, Philadelphia.
- Bloomfield, A., & Badler, N. (2003a). *Enhanced collision perception using tactile feedback*. Technical Report MS-CIS-03-45, University of Pennsylvania, Philadelphia.
- Bloomfield, A., & Badler, N. (2003b). *A low cost tactor suit for vibrotactile feedback*. Technical Report MS-CIS-03-44, University of Pennsylvania, Philadelphia.
- Bloomfield, A., & Badler, N. (2007). Collision awareness using vibrotactile arrays. *Proceedings of the IEEE Virtual Reality 2007 Conference*.
- Burdea, G. C. (1996). *Force and Touch Feedback for Virtual Reality*. New York: Wiley.
- Cheng, L.-T., Kazman, R., & Robinson, J. (1996). Vibrotactile feedback in delicate virtual reality operations. *Proceedings of the Fourth ACM International Conference on Multimedia*, 243–251.
- Cholewiak, R., & Collins, A. (2000). The generation of vibrotactile patterns on a linear array: Influences of body site, space, and time. *Perception & Psychophysics*, 62(6), 1220–1235.
- Cholewiak, R. W., & Wollowitz, M. (1992). The design of vibrotactile transducers. In I. R. Summers (Ed.), *Tactile aids for the hearing impaired* (pp. 57–82). London: Whurr Publishers Ltd.
- Craig, J. C., & Johnson, K. O. (2000). The two-point threshold: Not a measure of spatial resolution. *Current Directions in Psychological Science*, 9(1), 29–32.
- Dillman, G., & Thomas, C. (1992). *Kyusho-Jitsu: The Dillman method of pressure point fighting*. Reading, PA: George Dillman Karate International.
- Fletcher, R. (1996). Force transduction materials for human-technology interfaces. *IBM System Journal*, 35(3–4), 630–638.
- Kaczmarek, K. A., Webster, J. G., Bach-y-Rita, P., & Tompkins, W. J. (1991). Electrotactile and vibrotactile displays for sensory substitution systems. *IEEE Transactions on Biomedical Engineering*, 38(1), 1–16.
- Kirner, T. G., Kirner, C., Kawamoto, A. L. S., Cantão, J., Pinto, A., & Wazlawick, R. S. (2001). Development of a collaborative virtual environment for educational applications. *Proceedings of the Sixth International Conference on 3D Web Technology*, 61–68.
- Knapp, L. (2001). Lost? Go vest, young man. *Wired*. Available at <http://www.wired.com/news/gizmos/0,1452,44376,00.html>.
- Kolasinski, E. M. (1995). *Simulator sickness in virtual environments*. Tech. Rep. 1027, US Army Research Institute, Alexandria, VA.
- Meehan, M., Razzaque, S., Whitton, M. C., & Brooks, F. P., Jr. (2003). Effect of latency on presence in stressful virtual environments. *Proceedings of the IEEE Virtual Reality 2003 Conference*, 141–148.
- Rupert, A. H. (2000). An instrumentation solution for reducing spatial disorientation mishaps. *IEEE Engineering in Medicine and Biology Magazine*, 19(2), 71–80.
- Shinagawa, Y., Nakajima, J., Kunii, T. L., & Hara, K. (1997). Capturing and analyzing stability of human body motions using video cameras. *Computer Animation*, 48–57.
- Sun, X., Chen, C.-W., & Manjunath, B. S. (2002). Probabilistic motion parameter models for human activity recognition. *Proceedings of the 16th International Conference on Pattern Recognition*, 1, 443–446.
- Terry, C. M. (2006). The martial arts. *Physical Medicine and Rehabilitation Clinics of North America*, 17(3), 645–676.
- van Erp, J. B. F. (2000). Tactile information presentation: Navigation in virtual environments. *Proceedings of the First International Workshop on Haptic Human-Computer Interaction*, 165–173.
- Weinstein, S. (1968). Intensive and extensive aspects of tactile sensitivity as a function of body part, sex, and laterality. In D. R. Kenshalo (Ed.), *The Skin Senses*, chap. 10, 195–222. Springfield, IL: Charles C. Thomas.

Copyright of *Presence: Teleoperators & Virtual Environments* is the property of MIT Press and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.