

Virtual Peg-in-Hole Performance Using a 6-DOF Magnetic Levitation Haptic Device: Comparison with Real Forces and with Visual Guidance Alone

B. J. Unger,^{*} A. Nicolaidis,[‡] P. J. Berkelman,^{*†} A. Thompson,[‡]
S. Lederman,[§] R. L. Klatzky,[‡] and R. L. Hollis^{*}

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Abstract

We describe two experiments using three testbeds (real, virtual and vision-only) for comparison of user performance during 3-D peg-in-hole tasks. Tasks are performed using a six-degree-of-freedom (6-DOF) magnetic levitation haptic device. The experimental design allows a user to experience real and virtual forces using the same device. The first experiment compares real and virtual tasks. In the virtual task, a peg and hole are rendered haptically and visually. During the real task, a physical peg is attached to the underside of the haptic device. A hole in a plate attached to a force/torque sensor receives the peg. The second experiment compares a virtual haptic task to one performed using vision alone. Preliminary results indicate increased task time, more variation in force and position, and more failures occur with the virtual task than with the real task. More variation in force and position, and more failures occur with the vision-only task than with the virtual task. Users apply similar strategies for virtual and real tasks. Virtual haptic display, while worse than reality, contributes significantly to task performance when compared to vision alone.

1 Introduction

The use of haptic feedback for task performance in real and virtual environments has received considerable attention in recent years. Many haptic displays have been tested using various performance criteria.

The fidelity of a particular haptic display is often measured in terms of kinematic and dynamic design constraints such as force bandwidth and dynamic range [1] or frequency response and steady state accuracy [2]. Other tests have concentrated on the ability of a human

operator to perform specified tasks. Analysis of operator task performance has generally focused on binary failure/completion criteria, accuracy [3] or completion time analysis [4, 5].

While simple measurements of task performance demonstrate gains when haptic feedback is employed, they fail to delineate the underlying strategies used by the subject in attaining the goal. More sophisticated analysis, by Hannaford *et al.*, employed force/torque data collected throughout a task to provide a richer, quantifiable performance metric [6]. By examining force/torque data continuously during the procedure, a larger task can be broken into subtasks. This allows quantitative analysis of the effect of different parameters on each subtask. Identification of important subgoals, user force and position strategies, and the influence of device parameters may then provide guidance for improved interface design and further understanding of the psychophysics of haptics.

Vision plays a large role in the performance of many human tasks. Klatzky *et al.* contend that the haptic and visual systems have separate encoding pathways with the haptic pathway encoding for substance and the visual pathway for shape [7]. It seems possible that, by providing additional information, haptics with vision could improve task performance over vision alone. The extent to which task performance depends on vision and can be improved by haptic feedback is yet to be fully determined. Sato *et al.* actually found a larger range of forces was used when vision and touch were used together than when vision alone was used in task performance [8]. Others have found that the combination of haptic and visual modalities is more effective than vision alone [9, 10]. Richard and Coiffet determined that visual display of force feedback information was inferior to haptic display of the same data [3]. This seems to indicate that vision cannot substitute directly for the haptic channel. By comparing a task performed with haptics and vision to an identical one performed with vision alone, it may be possible to learn the relative values of haptic and visual feedback.

A simple square peg-in-hole placement task with hap-

^{*}The Robotics Institute, School of Computer Science.

[†]Now at John Hopkins University.

[‡] Department of Psychology, Human Computer Interaction Institute.

[§] Department of Psychology, Queen's University, Kingston, Ontario, Canada

tic and graphical feedback was selected for study. Such a task requires moderately complex movement and force application in 6 DOFs and has been previously used for task performance measurements [6]. Peg-in-hole task contact states and manipulation skills have also been studied with respect to automated [11] and human [12] assembly tasks. These studies may provide a point of reference for our goal of understanding human manipulation strategies.

To identify elements of task performance which are related to the environment in which the task is performed (real, virtual, or vision-only), it is important to control for differences in experimental setup. Ideally, the subject should not know which modality is being employed. The experiments we have designed use the same interfaces for all environments: real, virtual and vision-only. Subjects use the same manipulandum, connected to either a real or a virtual peg for all tasks. In each case, the peg can only be seen as a graphical representation on screen. This unique setup helps to ensure that discrepancies in task performance are due mainly to the differences between the real, virtual, and vision-only modalities.

The peg-in-hole task involves discrimination of point, edge, and face hard contacts during motion in 6 DOF. To compare virtual with real task performance, it is important that the simulation’s haptic feedback realistically represents this environment. Device limitations, such as maximum stiffness, position resolution and bandwidth, may result in noticeable deviations from the ideal haptic sensation. Whereas it is our purpose to examine the differences between a virtual haptic interface and a real one, if these differences are too great, users may adapt radically different strategies, making analysis difficult. By using a 6-DOF haptic device with excellent performance characteristics [13] we hope to eliminate this problem.

2 Magnetic Levitation Haptic Device

To effectively compare real, virtual, and vision-only tasks, it is necessary to use a rendering device capable of providing realistic haptic sensation. Ideally, such a device should have high position and force bandwidths, fine position resolution and high stiffness. In addition, 6 DOFs are necessary to emulate the forces and torques encountered in real 3-D peg-in-hole placement. The magnetic levitation haptic device used in this study, shown in Fig. 1, provides such a platform. The device is composed of a hemispheric actuator assembly, optical position sensors, electronics, and realtime computer.

The actuator consists of a hemispheric aluminum shell (flotor) which is enclosed within the stator’s fixed magnet assemblies. Six coils on the inner surface of the flotor provide actuator forces. The current in each coil interacts with strong magnetic fields of the enclosing




Figure 1: Magnetic levitation haptic device cut-away view of design.

magnets to produce six independent Lorentz forces, providing an arbitrary force/torque wrench on the flotor, and hence to the attached manipulandum and the user’s hand. Three LEDs on the flotor are observed by fixed optical sensors, providing realtime position and orientation information with resolutions of 5-10 μm depending on position in the workspace. Because of the low flotor mass and the essentially frictionless environment, high position bandwidth can be achieved (~ 125 Hz at ± 3 dB) [13]. Maximum stiffness is approximately 25 N/mm in translation and 50.0 Nm/rad in rotation [14]. 6-DOF motion of the handle has a range approximately that of comfortable fingertip motion with the wrist stationary (± 12 mm translation and $\pm 7^\circ$ rotation in all directions).

The magnetic levitation haptic device communicates with an SGI Indigo 2 workstation via 10 Mb Ethernet. For the experiments reported here, the virtual peg-in-hole environment was modeled by Baraff’s CoriolisTM software (see [15]).

3 Experimental Setup

Experiments were designed to quantitatively analyze the force and position strategies used in a peg-in-hole task. During the task the subject placed a square peg in a square hole. The peg and hole were represented graphically in 3D while the subject received haptic input from the task via a 6 DOF manipulandum. Three testbeds were designed: a real testbed in which the peg and hole were physical objects, a virtual testbed in which the subject encountered a graphically and haptically rendered peg and hole, and a vision-only testbed in which the subject could only see the peg and hole and had no haptic feedback.

To compare the subject’s force and position strategies

under these three conditions it is necessary to eliminate external influences as much as possible. Therefore, the manipulandum of the haptic device used to represent the virtual haptic peg is also used as the real haptic peg. It is also used as the input device for the vision-only condition. The subject is thus presented with a working environment which is nearly identical for all three modalities.




Figure 2: Virtual spring-damper coupling between simulation and haptic device

The virtual testbed uses a workstation performing peg-in-hole haptic and graphics rendering calculations. The peg and hole are displayed by the haptic device and on a computer screen respectively. The simulation and the haptic controller are coupled using a virtual spring and damper as shown in Fig. 2.

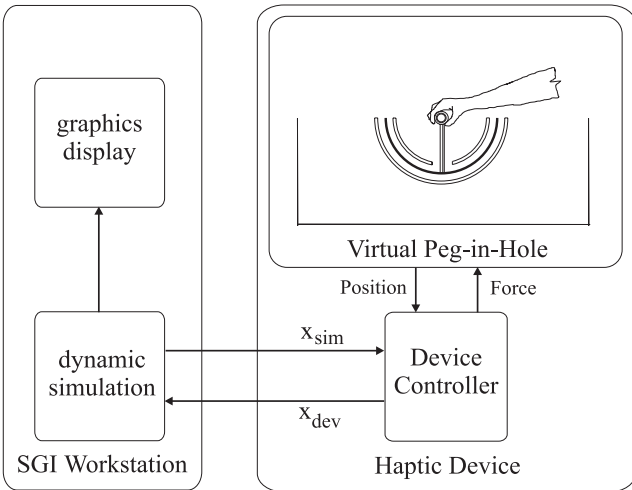


Figure 3: Experimental setup for virtual haptic peg-in-hole task.

Positions/orientations from the simulation (x_{sim}) and controller (x_{dev}) are exchanged as vectors between workstation and haptic device as shown in Fig. 3 and act as impedance control setpoints. Position error ($x_{sim} - x_{dev}$) and velocity feedback (v_{dev} or v_{sim}) provide the virtual spring-damper connection between the systems. The forces acting on the haptic device (f_{dev}) and present in the virtual representation f_{sim} are given by

$$f_{dev} = K_p (x_{sim} - x_{dev}) + K_v v_{dev}, \quad (1)$$

and

$$f_{sim} = \frac{K_{spring}}{K_{damp} v_{sim} + f_{other}} (x_{dev} - x_{sim}) + \quad (2)$$

where f_{other} are contact forces in the simulation. K_{spring} and K_{damp} are the gains of the virtual spring in the simulation and K_p and K_v are the gains for virtual spring of the haptic device. Gravity is cancelled by an additional feedforward term added to the z -axis forces in the simulation. The user's forces and torques required during task, along with position and orientation data are recorded at 100 Hz during the haptic device servo loop. Since the simulation calculates forces and positions by numerical integration of differential equations [15, 16] it runs slowly (30-50 Hz) compared to the haptic servo loop. The simulation update rate is also dependent upon the number of contacts occurring in any given time step. This can occasionally give rise to small latencies between haptics and graphics.

The vision-only test bed uses almost the same setup as the virtual haptic test bed. However, by turning off the amplifiers on the haptic device, force feedback is eliminated and the haptic device becomes essentially a position sensor. Position setpoints are still exchanged with the workstation simulation, allowing the subject to see the motion of the peg with respect to the hole. However, the subject no longer experiences any force feedback from the virtual objects. Simulation and haptic servo loops run at the same rates as in the virtual haptic test bed. The haptic servo loop continues to calculate and record forces, even though these forces are not experienced by the subject. The position/orientation of the flotor and the calculated forces/torques are recorded at 100 Hz by the haptic servo loop.

The real test bed uses a JR3* force/torque sensor mounted underneath the haptic device, within its enclosure and out of sight of the subject. A 12.75 cm square Delrin™ plastic plate containing a central 10.82 mm square hole is mounted on the JR3*. The entire assembly is placed directly beneath the center of the flotor. Ordinarily, the flotor is powered through wires connected at its lower pole. In the real test bed, however, the connector was replaced by a square Delrin™ peg as shown in Fig. 4. The square peg has a width of 10.72 mm allowing 0.1 mm clearance. The hole depth is approximately 10 mm. The subject sees the same graphical scene as that rendered for the virtual haptic task, now driven only by the haptic device position/orientation sensors. The forces/torques applied through the peg to the plate and hole are measured by the JR3* and recorded. The position/orientation of the flotor is recorded at 100 Hz by the device servo loop which no longer outputs forces/torques to the haptic display.

In the real, virtual haptic and vision-only test beds, the subject manipulates the peg using a T-shaped handle attached to the center inner surface of the flotor. The distance from the handle to the tip of the real peg is 19.2 cm. All three test bed tasks are designed to be nearly the same. By presenting the subject with the same graphical scene, differences in visual cues are avoided. By positioning the real peg beneath the flotor,

*JR3 Inc., Woodland, CA.




Figure 4: Experimental setup for real peg-in-hole task.

the subject is constrained to use similar hand positions for all three tasks. The hole position is similar for all test bed tasks, situated directly beneath the center of the flotor. Both the real peg and the plate containing the real hole are made using a plastic with a low coefficient of friction and friction is not modelled in the virtual task. (There were some unavoidable differences between the test beds, discussed in Sec. 6.)

4 Experimental Protocol

Two sets of experiments were conducted using the real, virtual and vision-only test beds. The first experiment examined subject performance during the peg-in-hole task in the real and virtual test bed environments. The second experiment compared subject performance during the same task in virtual and vision-only test bed environments.

During each experiment subjects are seated approximately 60 cm in front of a 19 inch SGI monitor displaying a graphical representation of a peg and a hole. The subject's viewpoint is above and to the left of the hole, providing a good view of the peg and the hole itself. The haptic device is placed so that the subject's outstretched arm makes a 45° angle with plane of the monitor as in Fig. 6.

The subject moves the peg to a predetermined start location indicated by a rectangular visual target to the right of the hole (see Fig. 5). On a signal from the examiner, the subject attempts to place the peg in the hole and data collection is initiated. Peg position and orientation, as well as force and torque data are recorded for a maximum of thirty seconds. The trial ends when the subject has successfully placed the peg in the hole and the examiner terminates the recording. Alternatively, if the subject is unsuccessful, recording terminates automatically at the end of thirty seconds.

Subjects were selected from a student subject pool. Only right-handed subjects were tested. Subjects were allowed a familiarization period with the apparatus prior to beginning trials. Each subject performed two blocks of trials. A block consisted of ten trials using a partic-




Figure 5: Graphical interface showing start target area and peg prior to placement in hole.

ular test bed (real, virtual, or vision-only). The trial modality performed first was counterbalanced to avoid training bias in the data. Subjects were not given any information regarding the modality they were using. Trials with a length of greater than 30 seconds were recorded as failures and not otherwise incorporated into the data.

In the first experiment, trials were performed using 9 subjects. Each subject performed twenty trials in a one hour time period. One block of trials was performed using the virtual peg-in-hole with haptic feedback and the other block of trials used the real peg-in-hole.

In the second experiment, trials were performed using 20 new subjects. Each subject performed twenty trials in a one hour time period. One block of trials was performed using the virtual peg-in-hole with haptic feedback and the other block of trials used only visual feedback.

For the purposes of analysis, the position of the peg was recorded in the fixed reference frame of the haptic device. In this right-handed frame of reference the positive z -axis is up while the positive x -axis is to the subject's right. Roll, pitch and yaw of the peg were also recorded.

Analysis of data was conducted on successful trials in both experiments. Unavoidable differences in peg and hole scales and offsets between real, virtual and vision-only conditions made descaling and offset subtraction necessary. Once all data was in the same metric, the start and end point of each trial was determined. The starting point was defined as the point where the peg's x axis position first passed the edge of the start target box. On successful trials, the first point where the peg position was 75% of the distance between the surface of the block and the bottom of the hole was considered the end of the trial. The interval from start to end was considered to be the length of a successful trial.



Figure 6: Subject using haptic device during peg-in-hole task.

5 Results

One of the assumptions underlying virtual reality research is that more accurate representation of physical objects provides a greater sense of virtual presence for the user. During a task, as more sensory channels are integrated into the experience of the virtual environment, we would expect to see improved task performance. Even with the current state of the art, simulated haptic feedback is not able to perfectly reproduce a real environment. Our experiments thus examine task performance in environments with differing levels of sensory input, ranging from excellent multi-modality to degraded multi-modality to single-modality.

We would expect that task performance would be at its best when manipulating a real tool. This is confirmed by the results of our first experiment, summarized in Table 1.

During the first experiment a total of 90 trials were recorded for the real task while 89 trials were recorded using the virtual haptic test bed. Subjects performed

the real peg-in-hole task both faster and more accurately than the virtual one. This can be seen from the shorter trial lengths and better success rates in the real environment. Terminal forces (forces during the last 1 second of a trial) applied by subjects in both real and virtual conditions were not significantly different in any axis. However, the variability of force application during any given trial, as measured by “within trial” terminal force standard deviation (σ), was significantly greater for the virtual haptic task.

The mean position data is difficult to interpret and may suggest differences in user positioning strategy between virtual haptic and real tasks. It could also reflect unrecognized deficiencies or biases in our experimental setup. Variability of position within trials was significantly higher for virtual tasks while variability of orientation was significant in two of the three axes.

The vision-only task provides sensory input of only one modality. If virtual haptic feedback provides useful additional information we would expect that task performance would improve when it was included. The results of our second experiment confirmed this (see Table 2).

We note that while trial length for virtual haptic and visual tasks is nearly identical for successful trials, success rates for the visual trials are significantly higher. It is also significant that success rates for the virtual component of the second experiment are similar to those in the same component of the first experiment. This indicates that the data are consistent and not subject dependent.

Mean force data is difficult to interpret for these trials. The calculated forces in the vision-only case are never displayed to the subjects. They may, as a result position the peg beneath the surface of the block without receiving visual feedback or being able to feel the resulting forces.

Mean position data is likewise suspect although a significant difference in the average y -axis position may indicate a difference in user positioning strategy. More interesting is the finding that position within vision-only trials varies significantly more than within virtual trials. Forces and torques within trials also tended to vary more significantly under vision-only conditions. This could also indicate that subjects found this modality more difficult, requiring the application of a wider range of forces. Orientation variation was, for the most part, not significantly different between trial modalities.

Interesting information can also be obtained by examining the position and force recordings directly. In Fig. 7 comparison of the positions of the peg, for representative trials, are seen side by side for the real and virtual haptic tasks. Differences in absolute position and scale are noticeable due to differences in the experimental setup, but similarities in strategy can be still be observed. For example, in looking at the simulated task x axis and z axis simultaneously, we note that the subject slides the peg along the surface (at a z -axis position of 2 mm) towards the hole which is located at an x -axis

	Virtual mean	Virtual between-subject σ	Real mean	Real between-subject σ	t test, critical value = 2.31, 2-tail
Duration of successful trials (s)	11.3	1.55	4.1	1.54	9.22
Success rate	0.84	0.12	1.00	0.00	-3.80
x -axis position (mm)	0.13	0.31	0.33	0.37	-1.23
y -axis position (mm)	-0.11	0.10	0.61	0.49	-4.84
z -axis position (mm)	3.10	0.36	2.42	0.16	5.15
Roll(degrees)	0.16	0.80	-0.27	0.30	1.99
Pitch (degrees)	0.23	0.84	-0.15	0.24	1.18
Yaw (degrees)	-0.33	0.79	-0.78	0.34	1.49
Terminal x -axis force (N)	-0.05	0.09	0.24	0.49	-1.73
Terminal y -axis force (N)	-0.16	0.14	-0.19	0.69	0.04
Terminal z -axis force (N)	-6.17	0.35	-6.73	0.45	0.61
Within-trial x position σ (mm)	0.90	0.31	0.59	0.08	3.81
Within-trial y position σ (mm)	0.38	0.16	0.21	0.05	6.72
Within-trial z position σ (mm)	0.77	0.31	0.39	0.12	6.70
Within-trial roll σ (degrees)	0.59	0.12	0.16	0.05	10.72
Within-trial pitch σ (degrees)	0.68	0.21	0.37	0.10	4.27
Within-trial yaw σ (degrees)	0.53	0.13	0.40	0.11	1.79
Within-trial x -axis force σ (N)	0.63	0.18	0.31	0.16	4.44
Within-trial y -axis force σ (N)	0.64	0.17	0.32	0.11	6.45
Within-trial terminal z -axis force σ (N)	1.82	0.21	1.23	2.70	-3.33

Table 1: Virtual haptic *vs.* real forces/torques and positions/orientations: means and average within-trial variability.

position of 0 mm. The y -axis deviations are small, indicating that a relatively two-dimensional path is used to reach the hole. Upon arriving at the hole, a series of lifts and drops are performed, detectable on the z axis as upward and downward deflections, until finally the peg is placed in the hole. A similar strategy is adopted by the subject for the real task, although within a shorter time frame and using smaller lifts. Initially, the peg is on the surface at approximately a z -axis position of -2 mm and x -axis position of 10 mm. The subject moves the peg towards the hole by sliding along the surface. The large deviation in the y -axis just as the peg enters the hole is likely due to momentary tipping of the peg resulting in motion in the y direction. As the peg aligns with the hole, this deviation is quickly corrected.

Examining the z -axis force recordings in Fig. 8, we can interpret the subject’s force strategies as they approach the hole. In both real and virtual haptic tasks we see that at (a) the peg is initially at rest. The peg’s weight and the weight of the subject’s hand resting on it account for most of the force. The force then decreases rapidly at (b) as the subject picks it up or slides it along the surface. Downward spikes (c) represent lifts, where the subject seeks to realign the peg, and upward spikes (d) represent taps, where the peg is dropped. The final spike (e) on the simulation data likely represents the weight of the subject’s hand as the peg is forced into the hole and comes to rest again at (f). From this virtual task tracing it is apparent that the subject repeatedly misaligned the peg and lifted it to clear it. The real task data reveals only a single momentary misalignment and

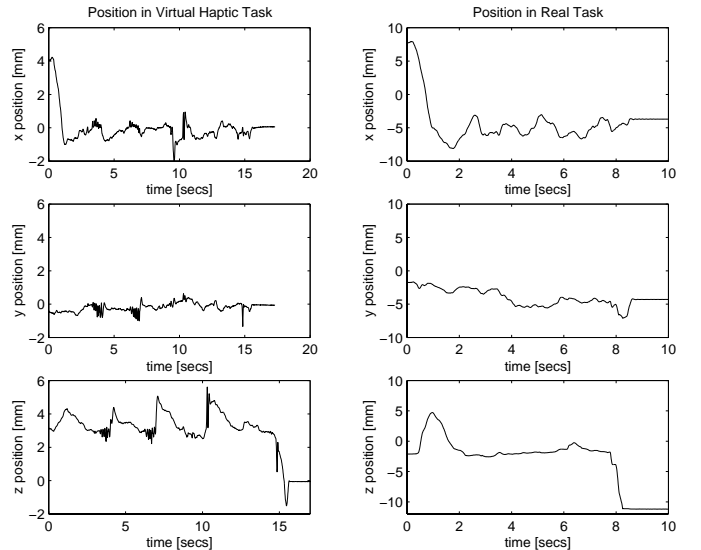


Figure 7: Comparison of peg position in real and virtual haptic tasks.

then success.

6 Discussion

Our findings indicate that task performance is best when using a real peg in a real hole. When a task is performed using the virtual environment, performance declines but is worst when the user experiences no haptic feedback.

	Virtual mean	Virtual between-subject σ	Vision-only mean	Vision-only between-subject σ	t test, critical value = 2.09, 2-tail
Duration of successful trials (s)	12.5	3.0	12.9	3.1	-0.35
Success rate	0.82	0.17	0.68	0.16	2.36
x -axis mean position (mm)	0.06	0.21	0.08	0.27	-0.43
y -axis mean position (mm)	-0.19	0.10	0.20	0.23	-8.34
z -axis mean position (mm)	3.36	0.22	2.22	0.53	9.42
Terminal x -axis force (N)	-.18	0.30	-.07	0.75	-1.79
Terminal y -axis force (N)	-.19	0.28	-.49	1.03	2.32
Terminal z -axis force (N)	-6.00	0.42	-5.65	1.26	0.95
Within-trial x position σ (mm)	0.81	0.16	0.95	0.20	-4.04
Within-trial y position σ (mm)	0.43	0.09	0.58	0.17	-3.71
Within-trial z position σ (mm)	0.84	0.13	1.67	0.59	-6.45
Within-trial roll σ (degrees)	0.52	0.14	0.64	0.22	-1.98
Within-trial pitch σ (degrees)	0.53	0.17	0.48	0.13	0.98
Within-trial yaw σ (degrees)	0.41	0.14	0.63	0.24	-3.89
Within-trial x -axis force σ (N)	0.53	0.16	0.75	0.33	-2.27
Within-trial y -axis force σ (N)	0.54	0.13	1.03	0.90	-2.09
Within-trial z -axis force σ (N)	1.49	0.47	1.70	0.98	-0.53
Within-trial x -axis torque σ (Nm)	0.06	0.02	0.31	0.16	-7.02
Within-trial y -axis torque σ (Nm)	0.06	0.03	0.35	0.26	-5.25
Within-trial z -axis torque σ (Nm)	0.02	0.01	0.23	0.24	-4.01

Table 2: Virtual haptic *vs.* vision-only forces/torques and positions/orientations: means and average within-trial variability.

Task completion times, task failure rates, and variability in positioning and applied forces confirm this.

Differences in real and virtual haptic task performance are likely due to disparities between the user’s expected sensation and the feedback delivered by the haptic system. A variety of shortcomings in the overall haptic system may account for these findings. By examining data recorded during the tasks, we can see the strategies used to overcome the simulation’s limitations. Repeated lifting of a misaligned peg and retreating from the hole to try again are some of the more obvious strategies adopted.

Differences in virtual haptic and vision-only task performance may imply that haptic feedback provides the user with a greater sense of realism than does visual feedback alone. Greater variability in positioning and force application during the vision-only task suggests that the combination of sensory modalities allows better control of a manipulated object possibly by providing more channels of information to the user.

The small degree of variation in orientation seen in all tests (never larger than 0.68°) may indicate that, for a 3D task, the additional constraint of a surface (real or perceived) encourages the subject to use fewer than the available 6 DOFs. Previous studies indicate that such constraints can improve task performance [17]. Subjects may voluntarily constrain themselves to about 3 DOFs to make the task easier to accomplish. This would imply that artificial constraints for such tasks as a square peg-in-hole could be added to assist a user with disabilities.

Finally, it is notable that, in spite of the difficulties encountered by subjects in performing the virtual haptic task, the techniques used to accomplish real and virtual tasks were similar. This suggests that, in spite of differences in the perceived haptic sensation, peg-in-hole placement strategy is essentially fixed. Careful quantification of such strategies could help guide the design of task specific haptic interfaces in the future.

The work reported here is of a preliminary nature, subject to technical limitations that can be fixed in future implementations. We should note that some differences between tasks reported here were unavoidable due to design constraints and an attempt to create the same subjective “feel.” For example, it was necessary for the virtual peg to have a size of 1.84 mm, vs. 10.72 mm in the real peg. Both pegs, however, had a clearance of about 0.1 mm. All other distances in the virtual and vision-only task were scaled down by a factor of four. During the virtual task, a minor high-frequency “buzz” was felt in the haptic device handle that was not felt in the real task. With the real task, the flotor has a lower rest position, placing the bottom of the hole 12 mm lower than in the virtual task. This is a small difference and is generally not noticed by subjects. Finally, there is the slow (30-50 Hz) update rate of haptic rendering vs. the ~ 1000 Hz haptic device servo rate, as well as communication latencies. These effects and others not yet identified lead to the measured differences in performance between the real and virtual tasks.

Despite the objections just mentioned, we have suc-

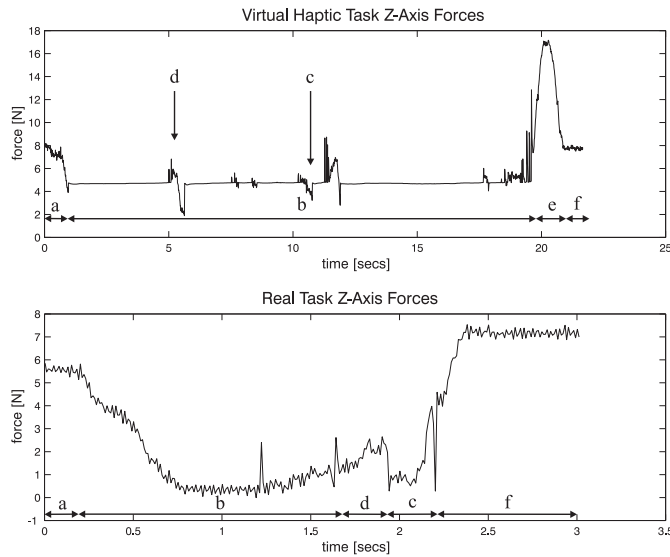


Figure 8: Comparison of z -axis forces during real and virtual haptic tasks: (a) peg at rest, (b) picked up, (c) lifting, (d) tapping, (e) hand forces peg down, (f) peg at rest in hole.

ceeded in obtaining a *quantitative* comparison between subjects performing real and virtual and vision-only versions of essentially the same 6-DOF task. This is in contrast to both engineering measurements (frequencies, bandwidths, resolutions, etc.) and to more subjective measurements (feels good, feels bad, not sure, etc.) that have been conducted previously.

7 Future Work

Further work can be done to bring the vision-only, virtual and real task setups into greater correspondence. Additionally, by moving the peg-in-hole haptic rendering algorithm from the workstation to the device controller, communication lags can be eliminated. Extensions could include task performance analysis under varying device and task parameters such as spatial resolution, number of DOFs, friction models and spatial tolerances. We want to quantitatively answer the question “how much reality can a haptic system provide”?

Acknowledgements

The authors would like to thank David Baraff for assistance with his CoriolisTM software, and the National Science Foundation for grants IRI-9420869 and IIS-9802191 which supported this work.

References

[1] R. Ellis, O. Ismaeil, and M. Lipsett, “Design and evaluation of a high-performance haptic interface,” *Robotica*, vol. 14, pp. 321–327, 1996.

[2] C. Hasser, “Tactile feedback with adaptive controller for a force-reflecting haptic display part 2: Improvements and evaluation,” in *Proceedings of the 1996 Fifteenth Southern Biomedical Engineering Conference*, pp. 530–533, 1996.

[3] P. Richard, A. Kheddar, and R. England, “Human performance evaluation of two handle haptic devices in a dextrous virtual telemanipulation task,” *Proceeding of the 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1543–1548, 1999.

[4] P. Buttolo, D. Kung, and B. Hannaford, “Manipulation in real, virtual and remote environments,” *Proceedings IEEE Conference on System, Man and Cybernetics*, vol. 5, pp. 4656–61, October 1995.

[5] R. Gupta, D. Whitney, and D. Zeltzer, “Prototyping and design for assembly analysis using multimodal virtual environments,” *Computer-Aided Design*, vol. 29, no. 8, pp. 585–597, 1997.

[6] B. Hannaford, L. Wood, D. McAfee, and H. Zak, “Performance evaluation of a six-axis generalized force-reflecting teleoperator,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 21, pp. 620–633, May/June 1991.

[7] R. L. Klatzky, S. Lederman, and C. Reed, “There’s more to touch than meets the eye: The salience of object attributes for haptics with and without vision,” *Journal of Experimental Psychology: General*, vol. 116, no. 4, pp. 356–369, 1987.

[8] D. Sato, S. Yamada, and M. Uchiyama, “Human skill analysis based on multisensory data,” *IEEE International Workshop on Robot and Human Communication*, pp. 278–283, 1997.

[9] P. Richard and P. Coiffet, “Human perceptual issues in virtual environments: Sensory substitution and information redundancy,” in *IEEE International Workshop on Robot and Human Communication*, pp. 301–306, 1995.

[10] L. Fabiani, G. Burdea, N. Langrana, and D. Gomez, “Human interface using the Rutgers Master II force feedback interface,” in *Proceedings of the IEEE Virtual Reality Annual International Symposium 1996*, pp. 54–59, 1996.

[11] J. Takamatsu, H. Kirnura, and K. Ikeuchi, “Classifying contact states for recognizing human assembly task,” in *Proceedings IEEE/SICE/RSJ International Conference on Multi-sensor Fusion and Integration for Intelligent Systems, 1999*, pp. 177–182, 1999.

[12] T. Yoshikawa and K. Yoshimoto, “Haptic simulation of assembly operation in virtual environment,” in *Proceedings of the ASME Dynamic Systems and Control Division*, vol. 69-2, pp. 1191–1198, ASME, 2000.

[13] P. Berkelman and R. Hollis, “Lorentz magnetic levitation for haptic interaction: Device design, performance, and integration with physical simulations,” *The International Journal of Robotics Research*, vol. 19, pp. 644–667, July 2000.

[14] P. J. Berkelman, *Tool-Based Haptic Interaction with Dynamic Physical Simulations using Lorentz Magnetic Levitation*. PhD thesis, Carnegie Mellon University, The Robotics Institute, 1999.

[15] D. Baraff, “Issues in computing contact forces for non-penetrating rigid bodies,” *Algorithmica*, no. 10, pp. 292–352, 1993.

[16] D. Baraff and A. Witkin, “Dynamic simulation of non-penetrating flexible bodies,” in *Computer Graphics (Proc. SIGGRAPH)*, vol. 26, (Chicago), pp. 303–308, Association for Computing Machinery, 1992.

[17] Y. Wang and C. MacKenzie, “The role of contextual haptic and visual constraints on object manipulation in virtual environments,” in *CHI Letters*, vol. 2, (The Hague, Amsterdam), pp. 532–539, April 2000.