Comparison of 3-D Haptic Peg-in-Hole Tasks in Real and Virtual Environments

B. J. Unger,* A. Nicolaidis,‡ P. J. Berkelman,*† A. Thompson,‡
R. L. Klatzky,‡ and R. L. Hollis*

Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Abstract

We describe an experimental arrangement for comparison of user performance during a real and a virtual 3-D peg-in-hole task. Tasks are performed using a unique six-degree-of-freedom (6-DOF) magnetic levitation haptic device. The arrangement allows a user to exert and experience real and virtual forces using the same 6-DOF device. During the virtual task, a peg and hole are rendered haptically, and visual feedback is provided through a graphical display. During the real task, a physical peg is attached to the underside of the haptic device. Using only real forces/torques, the peg is inserted into a hole in a plate attached to a force/torque sensor, while positions/orientations are measured by the haptic device. Positions/orientations and forces/torques are recorded for both modes. Preliminary results indicate increased task time, larger total forces and more failures occur with the virtual task. Recorded data reveal user strategies that are similar for both tasks. Quantitative analysis of the strategies employed should lead to identification of significant factors in haptic interface design and haptic rendering techniques.

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 performance criteria of various kinds.

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 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].

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, used by Hannaford *et al.*, employed force/torque data collected throughout a task to provide a richer, quantifiable performance metric [5]. By examining force/torque data continuously during the procedure, a larger task can be broken down 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.

A simple square peg-in-hole placement task with haptic 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 [5]. The combination of haptic and visual modalities has been previously studied and shown to be more effective than vision alone [6, 7]. Peg-in-hole task contact states and manipulation skills have also been studied with respect to automated assembly tasks [8]. 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 (virtual or real), it is important to control for differences in experimental setup. Ideally, the subject should not know which modality is being employed. The experiment we have designed uses the same graphical and haptic interfaces for both the real and virtual environments. The haptic user interface uses the same tool, connected to either a real or a virtual peg. In either case, the peg can only be seen as a graphical representation on screen. This unique

*The Robotics Institute, School of Computer Science.
†Now at John Hopkins University.
‡ Department of Psychology, Human Computer Interaction Institute.

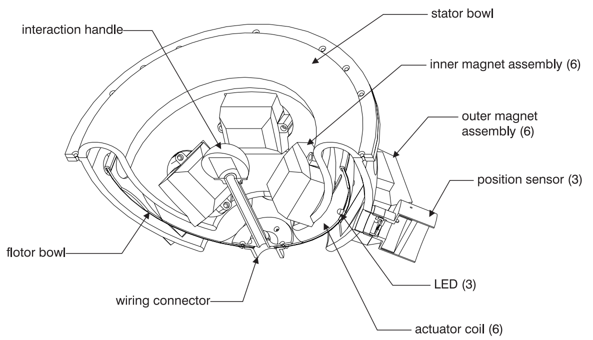

Figure 1: Magnetic levitation haptic device cut-away view of design.

setup helps to ensure that differences in task performance are due mainly to the differences between the real and virtual haptic interfaces.

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 [9] we hope to eliminate this problem.

2 Magnetic Levitation Haptic Device

To effectively compare real and virtual 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 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 ($\sim$125 Hz at $\pm$3 dB) [9]. Maximum stiffness is approximately 25 N/mm in translation and 50.0 Nm/rad in rotation [10]. 6-DOF motion of the handle has a range approximately that of comfortable fingertip motion with the wrist stationary ($\pm$12 mm translation and $\pm 7°$ rotation in all directions).

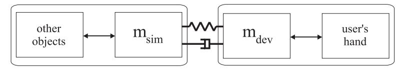

Figure 2: Virtual spring-damper coupling between simulation and haptic device

The magnetic levitation haptic device communicates with an SGI Indigo 2 workstation via 10 Mb Ethernet. For the experiment reported here, the virtual peg-in-hole environment was modeled by Baraff's Coriolis™ software (see [11]).

3 Experimental Setup

The task of putting a square peg in a square hole relies on discrimination of corner, edge and face contacts. In addition to haptic feedback, vision plays an important role in guiding the user's strategy.

The experiment consists of two 3D peg-in-hole tasks performed with 6 DOF haptic and visual feedback. The first task is a virtual one. The user is presented with a 3D graphical representation of a peg and hole, and places the peg in the hole while experiencing simulated haptic feedback. The second task requires the user to place a real peg into a real hole. This real task uses the same manipulandum attached to the peg and the same graphical representation of the peg and hole employed in the virtual task.

During the virtual task, the workstation performs peg-in-hole haptic rendering calculations which are displayed by the haptic device and graphics calculations which are displayed on the screen. The simulation and the haptic controller are coupled using a virtual spring and damper as shown in Fig. 2.

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 repre-

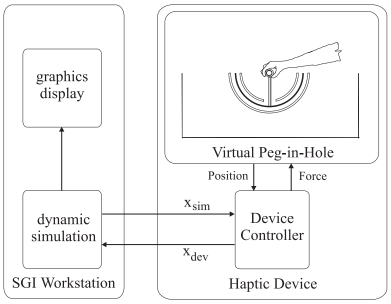

Figure 3: Experimental setup for virtual peg-in-hole task.

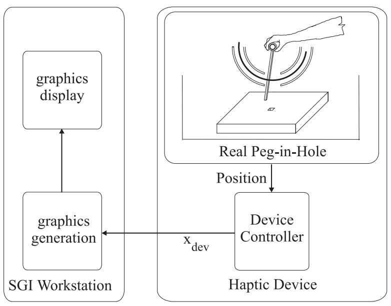

Figure 4: Experimental setup for real peg-in-hole task showing peg attached to flotor.

sentation f_{sim} are given by

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

and

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

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 calculated forces and torques required during servoing, 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 [12, 11] 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.

For the real task, a JR^3* force/torque sensor is 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 JR^3.

*JR^3 Inc., Woodland, CA.

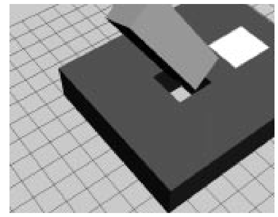

Figure 5: Graphical interface showing start target area and peg prior to placement in hole.

The entire assembly is placed directly beneath the center of the flotor. Ordinarily, the flotor is powered through wires connected at its lower pole. For the real task, 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 was approximately 10 mm. The subject sees the same graphics rendered for the virtual task, driven by the haptic device position/orientation sensors. The forces/torques applied by the subject to the peg are measured directly by the JR^3 and recorded. The position/orientation of the flotor are recorded at 100 Hz by the device servo loop which no longer outputs forces/torques.

In both the real and virtual tasks, the user manipulates the peg and receives haptic input from 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. The real and virtual tasks are designed to be nearly the same. By presenting the user with the same graphical scene, differences in visual cues are avoided. By positioning the real peg beneath the flotor bowl, the user is constrained to use similar hand positions for both real and virtual tasks. The hole position is similar for both 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 modeled in the virtual task. (There were some unavoidable differences between the real and virtual tasks, discussed in Sec. 6.)

4 Experimental Protocol

The experimental arrangement is designed to compare the force and positioning strategies used by subjects during the performance of a simple peg-in-hole task in real and virtual environments.

Each subject is seated approximately 60 cm in front of a 19 inch SGI monitor displaying a graphical representation of a peg and a hole (see Fig. 5). The haptic device is placed so that the subject's outstretched arm makes a 45° angle with plane of the monitor. After an initial familiarization period of not more than three minutes the trials begin.

The subject moves the peg to a predetermined start location to the right of the hole, as indicated by a visual target. On a signal from the examiner, the subject attempts to place the peg in the hole. Simultaneously, the examiner clicks the mouse and recording begins. 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, the recording terminates automatically at the end of thirty seconds.

Trials were performed using 9 subjects selected from a student subject pool. Only right-handed subjects were tested. Each subject performed twenty trials in a one hour time period. Ten of these trials were performed using the virtual peg-in-hole with haptic feedback and ten trials used the real peg-in-hole. The trial modality (virtual or real) performed first was counterbalanced to avoid training bias in the data. Subjects were not given information about which modality they were using. Trials which took longer than 30 seconds to complete were recorded as failures and not otherwise incorporated into the data.

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.

5 Results

The haptic senses can discriminate between very fine forces and positions, even while using tools. With the current state of the art in haptics, however, an approximation of reality is the best that can be achieved. We would therefore expect that task performance would be at its best when manipulating a real tool. This is confirmed by our study. Results are summarized in Table 1.

Parameter	Real Task	Sim. Task
Total Trials	90	89
Total Failures	0	14
Av. Failures/Subject	0%	15.8%
Av. Time/Trial [secs]	4.79	12.04
Shortest Trial [secs]	1.49	3.06
Longest Trial [secs]	15.28	26.07

Table 1: Time and success rates for tasks.

A total of 90 trials were recorded for the real task, while 89 trials were recorded for the simulation (due to loss of data on one trial). There were no failures during the real task but 14 failures (roughly 16 percent of attempts) were recorded for the simulation. Subjects took an average of 4.80 seconds to place the real peg in the hole and 12.04 seconds for the simulation. Substantial variation in trial length was observed for both real and simulated tasks. The minimum and maximum times for completion of the simulation were approximately double those for the real task.

Referring to Table 2, the mean position data indicates a bias about the center of the coordinate frame or in orientation of the peg, but the small σ for orientations reveals that only small changes were made during both real and simulated tasks; the exception being the yaw applied during the real task. This may be an artifact related to the positioning of the JR^3 force/torque sensor which was slightly misaligned with the coordinate frame axes.

Position	Real Mean	Real σ	Sim. Mean	Sim. σ
x-Axis [mm]	-1.77	2.27	0.16	0.88
y-Axis [mm]	-2.81	1.09	-0.10	0.36
z-Axis [mm]	-3.55	3.47	2.80	1.21
Roll [deg]	0.945	0.401	0.155	0.607
Pitch [deg]	0.516	0.728	0.258	0.670
Yaw [deg]	6.818	0.877	-0.424	0.510

Table 2: Mean and standard deviation of signed position and orientation of peg.

Interesting information can be obtained by examining the position and force recordings. In Fig. 6 comparison of the positions of the peg, for representative trials, are seen side by side for the real and simulated 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 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

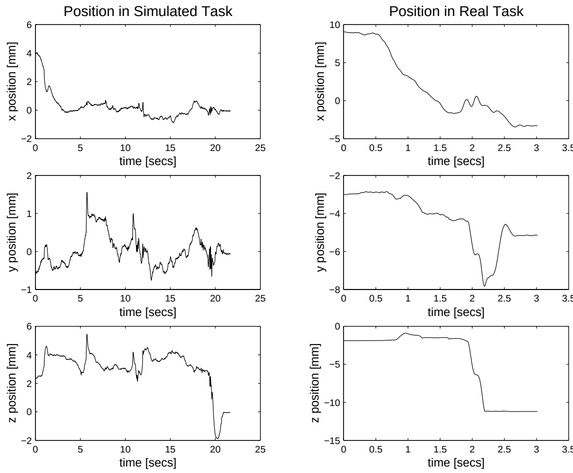

Figure 6: Comparison of peg position in real and simulated tasks.

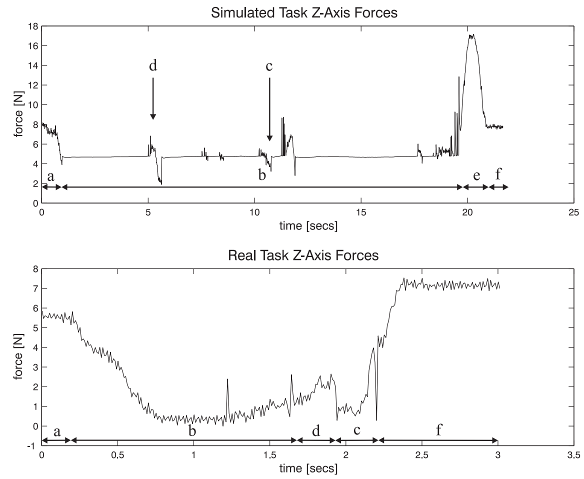

Figure 7: Comparison of z-axis forces during real and simulated tasks: (a) peg at rest, (b) picked up, (c) lifting, (d) jamming, (e) hand forces peg down, (f) peg at rest in hole.

peg enters the hole is likely due to momentary tipping of the peg resulting in motion of its center point in the y direction. As the peg aligns with the hole, this deviation is quickly corrected.

Observations of the force recordings are also revealing. A comparison of the z-axis forces reveals that the average applied force during a trial was 2.90 N for the real task, while it was 5.20 N for the simulation. This might be an indication that the subjects found the simulation more difficult and therefore it required greater application of force. Interestingly, while the average force is much greater for the simulation, the standard deviation of the average force required is almost the same (2.25 N real versus 2.35 N simulation). This would imply that the subjects required more overall force for the simulation, but were able to control the peg without large deviations from the norm.

Examining the z-axis force recordings in Fig. 7, we can interpret the subject's force strategies as they approach the hole. In both the real and simulated task 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 jammed peg, and upward spikes (d) represent taps, where the peg is dropped or jammed. 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 simulated task tracing it is apparent that the subject repeatedly jammed the peg and lifted it to clear it. The real task data reveals only a single momentary jam and then success.

6 Discussion

Our findings indicate that overall task performance is best when using a real peg in a real hole. Task completion time, overall average force and task failure rates all confirm this. These results are likely due to disparities between the subject'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 the position and force data recorded during the tasks, we can see the strategies used by the subjects to overcome the simulation's limitations. Repeated lifting of a jammed peg and retreating from the hole to try again are some of the more obvious strategies adopted.

It is also interesting to note similarities between the tasks. While the user applies more force during a simulation, the variation in the force is no greater than that in the real task. It is possible that the forces used in both tasks are small enough that the subject is able to apply the same degree of control over them. This result suggests that improving a user's ability to control a simulated tool could be achieved by scaling back the forces needed for performance.

The small degree of variation in orientation seen in both tests may indicate that, for a 3D task, the additional constraint of a surface encourages the subject to use fewer than the available 6 DOFs. Previous studies have indicated that such constraints can improve task performance [13]. It is possible that subjects voluntarily constrain themselves to about 3 DOFs to make the task easier to accomplish. This would imply that artificial constraints for tasks such as square peg-in-hole might 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 task, the techniques used to accomplish both tasks were similar. This suggests that, in spite of device limitations and differences in the perceived haptic sensation, the strategy for peg-in-hole placement 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 task were scaled down by a factor of four. During the simulated 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. All 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 succeeded in obtaining a *quantitative* comparison between subjects performing real and virtual 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 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"?

Aknowledgements

The authors would like to thank David Baraff for assistance with his Coriolis™ 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] B. Hannaford, L. Wood, D. McAffee, and H. Zak, "Performance evaluation of a six-axis generalized force-reflecting teleoperator," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 21, May/June 1991.

[6] 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.

[7] 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.

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

[9] 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.

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

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

[12] 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.

[13] 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.