

# Comparing Immersive Virtual Reality with Other Display Modes for Visualizing Complex 3D Geometry

David W. Mizell

Stephen P. Jones

*Boeing*

[David.Mizell@boeing.com](mailto:David.Mizell@boeing.com)

[Stephen.Jones@boeing.com](mailto:Stephen.Jones@boeing.com)

Mel Slater

Bernhard Spanlang

*University College London*

[m.slater@cs.ucl.ac.uk](mailto:m.slater@cs.ucl.ac.uk)

[b.spanlang@cs.ucl.ac.uk](mailto:b.spanlang@cs.ucl.ac.uk)

## Abstract

*This experimental research was aimed at determining whether or not immersive virtual reality (IVR) technology gives a user a measurable advantage over more conventional display methods when visualizing complex 3D geometry. Subjects were shown an abstract rod sculpture in a variety of display and display-control modes, and were tasked with assembling a physical replica of the sculpture they were visualizing. They were scored on the speed and accuracy with which they assembled their replica sculptures.*

*Head-tracked immersive VR was shown to have a statistically significant advantage over joystick-controlled display modes, especially in the case where the displayed sculpture was shown in super-scale, surrounding the subject.*

## 1. Introduction

Why use immersive virtual reality? IVR equipment is still rare and expensive. For viewing graphical datasets of any interesting size and geometric complexity, powerful, expensive graphics computers are also required. Standards, languages, APIs, tools and hardware interfaces are still not yet well established and widely accepted. Thus specialized technical people are also required for the installation and operation of the IVR system.

So what do we get for the money? What is the economic case to be made to the company CFO on behalf of IVR? Can we point to a realizable benefit? Can we cost-justify IVR? A technically sophisticated CFO will ask why we should buy an IVR system when, with the same amount of money, we could buy several quite powerful graphics workstations.

Hearing these kinds of questions asked in the industrial context was one of the motivations for this work. The economic questions prompted a research question: Can we demonstrate a measurable advantage IVR has over desktop displays, or other conventional human-computer interfaces?

It is clear that at least certain forms of VR offer a capability that simply doesn't exist for conventional, "flat screen" interfaces. The primary example is the quasi-physical

interaction made possible within a virtual environment by such devices as data gloves, with tracking sensors on the hand and other parts of the body. An IVR user can wear an animated graphical body, i.e., an "avatar," and reach toward, grasp and manipulate virtual objects. With the advent of haptics technology, it is even becoming possible to feel the solidity and weight of the virtual objects.

We decided to focus this inquiry entirely on the aspect of visualization. Granted that IVR offers unique forms of physical or quasi-physical interaction coupled with visualization, does IVR provide any advantage over conventional displays for visualization alone? We refined this question further, again guided by aerospace industry experience: Does IVR provide an advantage over conventional displays for visualizing and understanding complex collections of 3D geometry? Large aerospace manufacturers such as Boeing design their products using 3D CAD. The typical design review entails a group of engineers in a conference room, looking at the CAD geometry on a large, front-projection screen on the wall. They get lost a lot, or at least confused about what it is they are looking at. This doesn't happen to them inside a real airplane. Would IVR mimic their interactions with the real world closely enough that they wouldn't get confused about what they were looking at in a virtual environment, either?

## 2. The Experimental Approach

Our goal was to design an experiment where a human subject's ability to comprehend and generate an internal mental map of a complex 3D object could be measured. This would give us a way to quantitatively compare the visualization capabilities of one user interface against another.

The basic task for our experimental subjects consisted of making a real copy of a virtual object. We decided to build a set of abstract rod sculptures, in order to have a test object without any intrinsic meaning. The base of each sculpture was an 18" by 16" sheet of pegboard (1/8" sheets of masonite with a grid of 1/4" diameter holes located 1" apart) with plywood backing. We cut 8" to 18" pieces of aluminum rods that were 1/4" diameter and fit snugly into the pegboard holes. The rods were bent into various abstract shapes. The "sculptures" were created by choosing 5 or 6 of the bent rods and inserting them into selected pegboard holes. Each pegboard base had a pair of concentric equilateral triangles painted on it to assist in orientation.

Each sculpture so constructed was then modeled using a 3D solid modeling tool (see Figs. 1 and 2).




Figure 1. Example rod sculpture used in the experiments.




Figure 2. 3D solid model of the rod sculpture shown in Fig. 1

While the sculptures used and the visualization interfaces varied, the experimental paradigm was always the same: a subject would be shown a completed sculpture, using some visualization interface. They would also be given an empty pegboard base and a set of rods, held in a separate, 5" by

12" sheet of pegboard. Subjects would be instructed to use their pegboard base and rods to construct, as quickly but as accurately as possible, a physical replica of the completed sculpture they were being shown. The set of rods would contain the shapes needed to construct a duplicate of the displayed sculpture, plus a few that didn't belong. Subjects were able to look at the displayed, complete sculpture as often as they wanted to, as they constructed their replica. We recorded the time each subject took to complete their replica, starting the clock at the time we made the completed sculpture visible to the subject and stopping the clock when the subject told us he/she was finished. We scored the accuracy of the replica according to the selection and placement of the rods. Each of the following would be scored as an error:

- ?? Selecting the wrong rod
- ?? Inserting a rod in the wrong hole in the base
- ?? Placing a rod in the wrong angular orientation (we allowed 45 degrees of leeway here)
- ?? It was sometimes possible to place Rod A's horizontal section over Rod B's horizontal section when Rod B was actually supposed to be above Rod A, due to looseness of the pegboard holes. On the rare occasions this occurred, it was scored as an error (the other three errors were roughly equally distributed).

There is an aspect of the design of an experiment like this that can only be addressed subjectively: Does this experiment measure what we are trying to measure? We are forced to postulate that the relative ability to quickly and accurately reconstruct a virtually displayed rod sculpture will reflect the relative efficacy of the visualization interface for understanding the complex geometry being viewed. Our argument in favor of this approach is necessarily intuitive. To construct a replica sculpture, a subject must successfully

- ?? compare and match the abstract shapes of the rods
- ?? perceive the spatial relationships between the rods, how they are arranged on the board and which hole each is located in.

Basically, we are arguing that for a subject to quickly and accurately replicate this abstract tangle of bent rods, he/she must visualize it thoroughly and understand it well.

## 3. Related Work

Several experimental projects have aimed at comparing IVR to other user interfaces in one way or another. Much of Slater's work has been focused on how effectively the user interface contributes to a sense of "presence" [10, 11, 12, 13, 14, 15, 16, 18]. Some work has compared the efficacy of various picking and placing metaphors within the same virtual environment [1, 19]. Boyd compared the effectiveness of a set of control and movement metaphors within an immersive virtual environment [2].

For much of VR's history, experimental work aimed at directly comparing the capabilities of an immersive interface to a more conventional interface was inhibited by the low quality of early VR equipment and other confounding effects such as low-frame-rate tracking and rendering. More recently, however, technology has reached the point that experimentally comparing VR with other modes has become a topic worth pursuing. Chung compared head-tracked versus non-head-tracked modes for targeting radiotherapy treatment beams [4]. Pausch *et al* compared a head-tracked mode versus a non-head-tracked mode for searching a virtual room to determine whether or not a certain letter was drawn on any of the walls, ceiling or floor. Robertson *et al* later extended this work to "fishtank VR," using a workstation augmented with shutter glasses for stereo [8].

Slater *et al* used the game of three-dimensional chess to compare the effectiveness of IVR with data glove touch interaction with the chess pieces to a workstation screen with mouse interaction with the pieces [14].

Ruddle *et al* [9] and Slater *et al* [10] each compared IVR to other interfaces in terms of how well a given interface helped a subject navigate through a large building. Buelthoff *et al* have compared how quickly subjects can recognize a 3D shape in a virtual environment versus other interfaces [3].

## 4. Phases of the Experiment

### 4.1 Boeing, 1994

This experimental approach was first conceived and used by the first author and his colleagues at Boeing in the summer of 1994 [5]. Subjects were shown, in an order varying by subject, a physically present completed rod sculpture, or a virtual sculpture displayed in stereo on a Fakespace Boom, or on a 1280x1024 SGI workstation screen. Orientation of the image on the workstation screen was controlled by the subject by manipulating a 3" by 4" wooden board with a Polhemus tracking sensor mounted on it. The orientation of the sculpture image was slaved to the orientation of the board in the subject's hand.

Throughout this series of experiments, we have endeavored to minimize subjects' confusion caused by unfamiliarity with the control aspect of the user interface. When we used the Boom, its thumb buttons had no function; the display was only moved to change the point of view of the completed sculpture. We used the Polhemus on the wooden board for control of the image on the workstation screen so that the subject wouldn't have to learn how to use a mouse or some other device to change the point of view. We believed that slaving the workstation's sculpture orientation to the Polhemus/board interface was sufficiently intuitive that subjects wouldn't have to be trained to use it.

The experiment design we used in this phase was between-subject – subjects would use only one interface, and speed and accuracy comparisons would be made with other subjects who used other interfaces. This was out of necessity, not logic; the Boom had to travel to conferences during that time. We quickly became convinced by the wide variability of subjects in performing the sculpture-copying task that a within-subject comparison would be much more meaningful.

Results of this phase were fairly inconclusive. Viewing a physical completed sculpture sitting on the table next to where they were building their replica was far superior in speed and accuracy to either of the two computer graphics interfaces. No significant, unambiguous difference was found between the Boom and the workstation. Mostly what we learned was the difficulty of implementing a VR experiment of this type in such a way that the results aren't confounded by idiosyncrasies of the equipment. The Boom swung away when subjects let go of it, so they tended to build their replica with one hand. Similarly, subjects had to hold the Polhemus board in the same position once they found the point of view at which they wanted to view the completed sculpture. Some subjects seemed quite intimidated by the Boom.

### 4.2 UCL, 1999

In the summer of 1999 we set up another round of rod sculpture experiments. This was structured as a within-subject experiment. Each subject would replicate 3 sculptures, randomly selected, using each of 3 interfaces, in random order:

- ?? physical – a completed sculpture would be physically present, on a table to the subject's right
- ?? an SGI 1600x1100 pixel wide-aspect-ratio monitor would display the completed sculpture. A wooden joystick (Fig. 3) with a Polhemus FasTrack sensor mounted on top provided 3DOF control of the orientation of the screen image. The joystick was designed to have enough friction in the mechanism that it would stay in whatever orientation it was in when the subject let go of it.




Figure 3. Co-author Bernhard Spanlang demonstrates constructing a replica sculpture copied from the completed sculpture represented

on the workstation screen. Wooden 3DOF joystick below workstation controls orientation of the screen image.

?? a Sony Glasstron 800x600 color stereo HMD, with a Polhemus FasTrack sensor mounted on it. The Glasstron optics allow their transmissivity to be adjusted between 0 and 20%. We set the transmissivity at 20%. We wanted the subject to be able to look rapidly back and forth between the completed sculpture that he/she was seeing in the Glasstron and the replica he/she was assembling on a table in front of him/her, without having to remove the HMD or flip it up and down. We put a black cloth over the area where the virtual completed sculpture appeared, to eliminate background confusion, and shone a 500W lamp on the table where the subject was building the replica. This enabled the replica sculpture to be easily visible through the Glasstron optics (Fig. 4).



Figure 4. Constructing a physical replica of the completed sculpture seen virtually using the Sony Glasstron HMD.

As a fourth task, each subject would redo the first task, replicating the same sculpture using the same interface, so that we could detect any learning that might be taking place over the course of the experiment. None was found.

Each subject was given a visual acuity test before the experiment began, to make sure that each subject had at least average vision in both eyes. Following their experimental tasks, each subject was given a standard Spatial Aptitude Test (SAT) [17], in order to detect any relationship between ability to perform the sculpture-replication task and general ability to comprehend and mentally manipulate 3D geometry.

36 subjects were used, one for each permutation of sculpture order times each permutation of interface order.

Results once again showed a significant superiority of the physical representation over the two computer graphics representations, but no statistically significant difference between the workstation/joystick and the Glasstron/head tracker interfaces. The results again seemed confounded by the idiosyncrasies of the equipment. The SGI display had much higher resolution than the Glasstron (we should have

thought of this and set it at 800x600). Ergonomics of the Glasstron HMD were a factor. It frequently slipped down, and was difficult to adjust to a comfortable, stable placement on the head. Subjects didn't seem to understand the effect of head tracking. They perceived that moving their head would change the image in the HMD, but tended to make unnatural cocking and tilting movements of the head, unlike what people do when examining a real object on a table from different points of view.

### 4.3 UCL, 2000

In May and June 2000 we ran another series of rod sculpture experiments. This time we used UCL's CAVE, which is equipped with four projections surfaces, three walls and the floor. We placed a low chair in the middle of the CAVE with a low table in front of it. The subject built the replica on the table. The 3DOF joystick was mounted on a separate, small table at the right hand of the subject. The completed sculpture was shown to the subject on the walls and floor of the CAVE. We decided that the physical representation of the completed sculpture had served its purpose as an experimental control in previous phases and did not use it in this phase.

The big advantage of using the CAVE was that several user interface characteristics could each be changed in the CAVE, independently of each other:

- ?? whether the completed sculpture was displayed in mono or stereo vision. The subject wore Crystal Eyes shutter glasses throughout the series of four tasks they performed, which were operating for two of the tasks and not for the other two.
- ?? head-tracked vs. joystick – the UCL CAVE has an InterSense IS-900 acoustic-inertial hybrid 6DOF tracking system. One tracking sensor is mounted on the brow frame of the Crystal Eyes, the other on a hand-held wand usable as a pointing or selecting device in the CAVE. We mounted the wand on top of the joystick, and either used the head tracker or the joystick, never the two together.
- ?? normal scale versus super-scale – a new variation was added for this phase of the experiments. The subject was either shown the completed sculpture on the right-hand wall and the area of the floor on his/her right, as in earlier phases, or was shown the sculpture increased in size to a much larger scale, roughly ten times natural size, with the subject appearing to be seated in the center of the sculpture.

We designed and modeled four new sculptures, which would fit within the display limits of the UCL CAVE's screens even in super-scale. Subjects always built replicas of the four sculptures in the same order, but the interface modes were permuted, with each subject using mono for two and stereo for the other two, head-tracking for two and joystick viewpoint control for the other two, and normal-

size for two and super-scale for two, in a Latin Square experiment design for 24 subjects.

The subjects were shown normal-scale completed sculptures on their right, similarly as earlier phases (Fig. 5). For half of such tasks, the head tracker would change the subject's view of the sculpture, and for the other half the orientation of the sculpture was slaved to the orientation of the joystick. These modes corresponded fairly closely to the HMD-vs.-workstation comparisons we made in previous phases, except that stereo vs. mono was now varied independently.



Figure 5. Subject in UCL CAVE, normal-scale/stereo/head-tracked mode (joystick is not in use and is tilted out of the way).



Figure 6. Subject in CAVE, super-scale/mono/head-tracked mode. Completed sculpture is visible on all display surfaces of the CAVE.

When the subject saw a sculpture in super-scale and head-mounted modes, it was displayed on the walls and floor around the subject, and subjects just had to turn their heads left and right to see it all (Fig. 6). In super-scale/joystick mode (independently of stereo vs. mono), only the right-hand wall was active (Fig. 7). The subject was instructed that the sculpture was super-scale and all around them, but was only visible on the right-hand wall, the "magic window." The orientation of the sculpture was slaved to the joystick, as usual, with the pivot axis located in the center of the virtual sculpture where the subject was sitting, so to see other parts of the sculpture the subject was instructed to rotate the joystick to make the sculpture rotate past the right-hand display wall.



Figure 7. Subject using super-scale/mono/joystick mode. This subject is in the last group of 22 who used the "camera-moves-the-world" mode, with the mockup camera mounted on top of the joystick.

The experiments were conducted in early June 2000. We finished 24 subjects and had more volunteers, so we began repeating the tasks in the 24-subject experiment design with more subjects. For this additional group of subjects, 22 in all, we changed one aspect of the user interface, to see if it would make a difference in the results. This time, when a subject was viewing a sculpture in super-scale/joystick mode, any orientation change in the joystick, instead of causing the same change in the virtual sculpture's orientation, would produce the *opposite* change. Twisting the joystick to the left would rotate the sculpture to the right, and so on. We called this mode "camera-moves-the-world" (perhaps "joystick moves the camera" would have been more accurate) and the earlier mode "joystick-moves-the-world." To reinforce the metaphor, we built a mockup camera and mounted it on top of the joystick (Fig. 7). Once again, only the right-hand wall of the CAVE was active, and subjects were instructed to think of it as the viewfinder or monitor of the joystick-camera.

#### 4.4 Results of the UCL 2000 CAVE Experiments

We were more confident in the CAVE experiments than in previous phases. Although we were not literally comparing IVR with a workstation, this setup had several aspects that made comparisons more fair:

- ?? The subject was always using the same display, with the same resolution and color characteristics.
- ?? The field of view of the Crystal Eyes shutter glasses was almost exactly one CAVE wall's worth. Thus, how much of a super-scale sculpture one sees at any instant in head-tracked mode is the same amount of image information as one sees on the right-hand wall in joystick mode.
- ?? We could vary stereo versus mono vision without removing the Crystal Eyes from the subject, so the physical characteristics of the display interface were the same throughout the experimental tasks. It turned out that some subjects never noticed whether they were seeing mono or stereo.

The experimental design we used enabled comparisons to be made within-subject for single modes (e.g. stereo vs.

mono) or for any pair of modes (e.g. head-tracked/stereo vs. joystick/mono), but triplets of modes (e.g. head-tracked/stereo/super-scale vs. joystick/mono/normal-scale) had to be compared between subjects. We didn't anticipate the triplet comparisons to be significant and our evidence supports this.

As one might expect, the super-scale representation caused the subject to take longer and make more errors than the normal-scale virtual sculpture, independently of the other parameters. For the joystick mode, on average the subjects took 75% longer with the super-scale representation and committed eight times as many errors. For the head-tracked mode, the subjects took nearly 50% longer with the super-scale representation and committed twice as many errors. For the normal-scale sculpture, they could see it all at once, at a similar size to their replica. In the super-scale case, they were building a scale model of something they saw, surrounding them. It wasn't as easy to compare the shapes of sculpture pieces. Errors were slightly higher for super-scale/HMD than for normal-scale/HMD, probably because the tables and chair blocked the view of part of the floor, and subjects couldn't always tell which hole a rod was placed in, in super-scale mode.

We found head-tracked mode to be consistently superior in time and error rate to joystick mode. The difference in completion times is statistically significant in the case of both super-scale and normal-scale representation of the completed sculpture, but the larger difference occurred in the case of super-scale. For the super-scale representation subjects took, on the average, over 40% longer and committed three times as many errors with the joystick mode. For the normal-scale representation subjects took, on average, 20% longer with the joystick mode. There was no significant difference in error rate for normal-scale.

This brings us to another interesting result. We never detected any statistically significant difference between stereo vision and mono. Normal-scale virtual sculptures were presented at an arm's-length distance, and some pieces of the super-scale sculptures came within similar distances to the subject, so stereopsis was definitely present. Experiment monitors outside the CAVE, not wearing Crystal Eyes, routinely saw separated double images in stereo mode. The explanation may be in the other stereo cues the brain uses besides stereopsis. The CAVE's InterSense tracker is very fast and smooth. Motion parallax may have been a sufficient cue for the subject to perceive 3D.

Results for the last 22 subjects were generally similar to the first 24. Joystick/camera mode took an average of 60% longer than head-tracked mode in the super-scale case, but there was no difference in error rate. Also, the camera-moves-world/super-scale mode took about 20% longer on

average than joystick-moves-world/super-scale had for the earlier group of 24 subjects, and the error rate was three times higher.

The SAT was found to be a significant predictor of completion time, but the differences between interfaces that we found were independent of the SAT scores.

## 5. Conclusions

These experimental results probably reinforce the intuition of many VR researchers in the sense that the biggest advantage between immersive VR and the modes representing the "flat-screen" workstation occurred when the completed sculpture was presented in super-scale. That indicates that immersion helps the most when the *data* is in some sense immersive, when the natural way to visualize a set of geometry is to be surrounded by it. Aircraft CAD geometry can serve as a canonical example.

## 6. Future Work

One intriguing thing occurred when the authors were preparing the hardware and software for the 1999 round of experiments at UCL. A bug in the VR software caused the system frame rate to suddenly drop from the usual 25+ frames/second to around 4 frames/second. One of the authors was looking at a virtual sculpture in Glasstron/HMD mode, thinking how good it looked and how low-latency the system was, compared to older VR hardware, when suddenly the bug kicked in and the frame rate dropped to 4 frames/second. The author found himself lost in a sculpture he himself had designed. It made us wonder if there is a frame rate threshold below which people's ability to visualize complex geometry is significantly impaired. Display frame rate would be straightforward to experimentally vary.

One question we wrestle with is the somewhat "mixed reality" flavor of this experiment. While previous work in comparing IVR to other interfaces fairly cleanly separated the types of interfaces (see [14]), in this experiment we have subjects rapidly switching back and forth between their view of a virtual object, and their view of the physical replica they are building on the table in front of them. This was chosen as an effective way to measure the amount by which the interface was helping them visualize complex geometry, but it raises the question of whether or not an experiment could be designed for addressing this question without forcing subjects to rapidly and repeatedly switch back and forth between virtual and real.

## 7. Acknowledgments

The authors wish to thank the UK's Engineering and Physical Sciences Research Council (EPSRC) and the Boeing Company's Mathematics and Computing

Technology organization, particularly Drs. Ken Neves and Al Erisman, for their financial support of this research.

Special thanks go to Mr. Bill Ahern of UCL's Estates and Facilities Carpentry Shop, for his design and construction of the tables and the high-friction joystick used in the UCL experiments, and to David Swapp, UCL's CAVE manager, for his help with this research.

## 8. References

- [1] K. Arthur, K. Booth and C. Ware, "Evaluating 3D Task Performance for Fish Tank Virtual Worlds," *ACM Transactions on Information Systems*, Vol. 11 No. 3, July 1993, pp. 239-265.
- [2] C. Boyd, "Does immersion make a virtual environment more usable?," CHI 97 Conference Companion, 325-326.
- [3] Bülthoff, H.H., M.O. Ernst, F.N. Newell and B.S. Tjan, "Visual and haptic recognition of objects: Effects of viewpoint", *Investigative Ophthalmology & Visual Science* 40, 398 (1999).
- [4] J. Chung, "A Comparison of Head-tracked and Non-head-tracked Steering Modes in Targeting of Radiotherapy Treatment Beams", in *Proceedings, 1992 Symposium on Interactive 3D Graphics*, April 1992, pp. 193-196.
- [5] D. Mizell, S. Jones, P. Jackson, D. Pickett, "Is VR better than a workstation? A report on human performance experiments in progress", in M. Goebel, ed., *Virtual Environments*, Springer, 1995, ISBN 3-211-82737-4, pp. 1-7.
- [6] R. Pausch, D. Proffitt, and G. Williams, "Quantifying immersion in virtual reality", *SIGGRAPH '97*.
- [7] R. Pausch, M. Schelford, D. A. Proffitt, "A User Study Comparing Head-Mounted and Stationary Displays", *IEEE Symposium on Research Frontiers in Virtual Reality*, 1993, pp. 41-45.
- [8] G. Robertson, M. Czerwinski and M. van Dantzich, "Immersion in Desktop Virtual Reality", *Proceedings, UIST 97*.
- [9] Ruddle, R., Payne, S. and D. Jones, "Navigating Large-Scale Environments: What Differences Occur Between Helmet-Mounted and Desk-Top Displays?" *Presence: Teleoperators and Virtual Environments*, Vol. 8 No. 2, April 1999, pp. 157-168.
- [10] M. Slater, C. Alberto, and M. Usoh, "In the Building or Through the Window? An Experimental Comparison of Immersive and Non-Immersive Walkthroughs", *Proceedings, Virtual Reality Environments in Architecture and Design*, 1994.
- [11] Slater, M. and Usoh, M. "Presence in Immersive Virtual Environments", *Proceedings of the IEEE Virtual Reality Annual International Symposium*, Seattle, WA, September, 1993, 90-96.
- [12] Slater, M. and Usoh, M. "Representation Systems, Perceptual Position and Presence in Virtual Environments", *Presence: Teleoperators and Virtual Environments*, 2.3 MIT Press, 1994, pp. 221-234.
- [13] Slater, M., Usoh, M. and Steed, A. (1994) "Depth of Presence in Virtual Environments", *Presence: Teleoperators and Virtual Environments*, 3.2, 1994, pp. 130-144.
- [14] M. Slater, V. Linakis, M. Usoh, R. Kooper, "Immersion, Presence and Performance in Virtual Environments: An Experiment with Tri-Dimensional Chess", *ACM Virtual Reality Software and Technology (VRST)*, 1996, Mark Green (ed.), ISBN: 0-89791-825-8, pp. 163-172.
- [15] M. Slater, A. Steed, J. McCarthy, F. Maringelli, "The Influence of Body Movement on Subjective Presence in Virtual Environments", *Human Factors*, 1998, (in press).
- [16] M. Slater and A. Steed, "A Virtual Presence Counter", *Presence: Teleoperators and Virtual Environments* 9.5 (October 2000), in press.
- [17] P. Smith and C. Whetton, *General Ability Tests (User's Guide)*, The National Foundation for Educational Research, ASE, 1998.
- [18] S. Uno and M. Slater, "The Sensitivity of Presence to Collision Response", *Proceedings of the IEEE Virtual Reality Annual International Symposium (VRAIS)*, Albuquerque, New Mexico, USA, April, 1997.
- [19] S. Zahi and P. Milgram, "Human Performance Evaluation of Manipulation Schemes in Virtual Environments" *Proceedings, IEEE Virtual Reality Annual International Symposium*, September 1993, pp. 155-161.