

Effects of Handling Real Objects and Avatar Fidelity On Cognitive Task Performance in Virtual Environments

Benjamin Lok
University of North Carolina at
Charlotte
bclok@cs.uncc.edu

Samir Naik
Disney Corporation
Samir.D.Naik@disney.com

Mary Whitton,
Frederick P. Brooks Jr.
University of North Carolina at
Chapel Hill
[whitton, brooks]@cs.unc.edu

Abstract

Immersive virtual environments (VEs) provide participants with computer-generated environments filled with virtual objects to assist in learning, training, and practicing dangerous and/or expensive tasks. But for certain tasks, does having every object being virtual inhibit the interactivity? Further, does the virtual object's visual fidelity affect performance? Overall VE effectiveness may be reduced if users spend most of their time and cognitive capacity learning how to interact and adapting to interacting with a purely virtual environment.

We investigated how handling real objects and how self-avatar visual fidelity affects performance on a spatial cognitive task in an immersive VE. We compared participants' performance on a block arrangement task in both a real-space environment and several virtual and hybrid environments. The results showed that manipulating real objects in a VE brings task performance closer to that of real space, compared to manipulating virtual objects.

1. Introduction

Conducting design evaluation and assembly feasibility tasks in immersive virtual environments (VEs) enables designers to evaluate and validate multiple alternative designs more quickly and cheaply than if mock-ups are built and more thoroughly than can be done from drawings. Design review has become one of the major productive applications of VEs [1].

The ideal VE system would have the participant fully believe he was actually performing a task. In the assembly verification example, parts and tools would have mass, feel and look real, and handle appropriately. The participant would naturally interact with the virtual world, and in turn, the virtual objects would respond to the participant's action appropriately [2].

Obviously, current VEs are far from that ideal system. Indeed, not interacting with every object as if it

were real has distinct advantages. In current VEs, almost all objects in the environment are virtual, but both assembly and servicing are hands-on tasks, and the principal drawback of virtual models — that there is nothing there to feel, nothing to give manual affordances, and nothing to constrain motions — is a serious one for these applications. Imagine trying to simulate a task as basic as unscrewing an oil filter from an engine in a VE!

Interacting with purely virtual objects limits the types of feedback the system can provide. In particular, the system's lack of constraints and haptics reduces the naturalness of interactions between real objects in the VE (including the participant) and virtual objects. Further, the VE representations of real objects (*real-object avatars*) are approximations and not necessarily visually faithful to the object itself. This paper reports the results of an investigation of the impact of these factors on task performance in a spatial cognitive task.

In this work, we extend the definition of an *avatar* to include a virtual representation of any real object, not just the participant. The *real-object avatar* is registered with the real object, and ideally, it is identical in look, form, and function to the real object. We use *self-avatar* to refer specifically to a participant's virtual representation.

We feel a *hybrid environment* system, one that incorporates representations of dynamic real objects into the VE, would assist in providing natural interactivity and visually-faithful self-avatars. These features should improve task performance, and in turn, VE effectiveness.

We believe spatial cognitive tasks, common in simulation and training VEs, would benefit from incorporating real objects. These tasks require problem solving while manipulating objects and maintaining mental model of relationships among them. The study we report on here is based on such a task.

2. Previous Work

The participant is represented within the VE with a self-avatar, either from a library of representations, a generic self-avatar, or no self-avatar. A survey of VE

research shows the most common approach is a generic self-avatar – literally, one size fits all [1]. Self-avatars are typically stylized human models, such as those in commercial packages. Although these models contain a substantial amount of detail, they do not visually match each specific participant’s unique appearance.

Researchers believe that providing generic self-avatars improves sense-of-presence over providing no self-avatar [3]. However, they hypothesize that the visual misrepresentation of self would reduce how much a participant believed he was “in” the virtual world, his *sense-of-presence*. Usoh concludes, “Substantial potential presence gains can be had from tracking all limbs and customizing [self-]avatar appearance [4].” If self-avatar visual fidelity might affect sense-of-presence, might it also affect task performance?

In general, VE systems attach extra trackers to the participant for sensing changing positions to drive an articulated self-avatar model. Matching the virtual look to the physical reality is difficult to do dynamically, though static-textured, personalized self-avatars are available in commercial systems such as AvatarMe [5].

A participant would ideally interact with the VE in the same way as he would in a real world situation, i.e. using his hands, body, and tools to manipulate objects in the environment. As a step towards this goal, some VE systems provide tracked, instrumented real objects as input devices. Common devices include articulated gloves with gesture recognition or buttons (Immersion’s Cyberglove), tracked mice (Ascension Technology’s 6D Mouse), or tracked joysticks (Fakespace’s NeoWand).

Another approach is to engineer a device for a specific type of interaction to improve interaction affordance. For example, tracking a toy spider registered with a virtual spider [6] or augmenting a doll’s head with sliding rods to enable doctors to more naturally visualize MRI data [7] enhances the VE interaction. However, this specialized engineering can be time-consuming and the results are often usable for only a particular task.

Several studies have been done on VE interaction, and [8] and [9] provide good summaries.

3. User Study

The study reported here was part of a larger examination of the effects of incorporating real objects into VEs. In the context of cognitive tasks, we asked the following questions:

- Does interacting with real objects improve task performance?
- Does seeing a visually faithful self-avatar improve task performance?

To investigate these questions, we employed a hybrid virtual environment system that can incorporate dynamic real objects into a VE. The system uses multiple cameras

to generate virtual representations of real objects at interactive rates [10].

The ability to include real objects in the VE allowed us to investigate how performance on cognitive tasks, i.e. time to complete, is affected by interacting with real versus virtual objects.

Video capture of real object appearance also has another potential advantage — enhanced visual realism. This feature enables the system to render a visually faithful self-avatar, which in turn allows us to investigate whether a visually faithful self-avatar, as opposed to a generic self-avatar, improves task performance. The results will provide insight into the need to invest the effort to use high visual fidelity self-avatars.

3.1. Task Description

We sought to abstract tasks common to VE design and training applications. Through surveying production VEs [1], we noted that a substantial number involved spatial cognitive manual tasks. We wanted to use a task that depended on cognition and manipulation, instead of participant dexterity or reaction speed, because of participant physical variability and given the current state of VE technology and applications.

We conducted a user study on a block arrangement task. We compared the participant’s task performance in three conditions: a purely virtual system and two hybrid systems that differed in level of visual fidelity. In all three cases, we used real-space task performance as a baseline.




Figure 1 – Three-inch wooden blocks were manipulated to match a target pattern.

The task we designed is similar to, and based on, the block design portion of the Wechsler Adult Intelligence Scale (WAIS). Developed in 1939, the Wechsler Adult Intelligence Scale is a test widely used to measure IQ

[11]. The block-design component measures reasoning, problem solving, and spatial visualization.

In the WAIS test, participants manipulate cubes to match target patterns. As the WAIS test is copyrighted, we modified the task to still require cognitive and problem solving skills while focusing on the interaction. We increased the size of the blocks from one-inch to three-inch cubes, as shown in Figure 1. The one-inch cubes would be difficult to manipulate in the purely virtual condition, and in the hybrid conditions, reconstruction errors in the relatively low-resolution real-object avatars of the cubes could hamper interaction.

Participants manipulated four or nine identical wooden blocks to make the top face of the blocks match a target pattern. The cubes were identically painted with the six possible quadrant-divided white-blue patterns.

There were two sizes of target patterns, *small* four-block patterns in a two-by-two arrangement, and *large* nine-block patterns in a three-by-three arrangement.

3.2. Study Design and Methods

The user study was a between-subjects design. Each participant performed the pattern matching task in a real space environment (RSE), and then in one of the three VE conditions. The independent variables were the VE interaction modality (real or virtual blocks) and the VE self-avatar visual fidelity (generic or visually faithful).




Figure 2 – Each participant conducted the task in the RSE and then in one of three VEs.

The three VE conditions were:

- Virtual objects with a generic self-avatar (*purely virtual environment - PVE*)
- Real objects with a generic self-avatar (*hybrid environment - HE*)
- Real objects with a visually faithful self-avatar (*visually-faithful hybrid environment – VFHE*)

The participants were randomly assigned to one of the three groups, 1) RSE then PVE, 2) RSE then HE, or 3) RSE then VFHE (Figure 2).

The target patterns were of a medium difficulty (determined through pilot testing). Our goal was to use target patterns that were not so cognitively easy as to be manual dexterity tests, nor so difficult that participant spatial ability dominated the interaction. All participants saw the same twenty patterns (six real space practice, six real space timed, four VE practice, four VE timed), the order of the patterns that each participant saw was unique.

Real Space Environment (RSE). The participant sat at a desk (Figure 3) with the nine wooden blocks inside a rectangular enclosure. The side facing the participant was open and the enclosure was draped with a dark cloth. Two lights lit the inside of the enclosure. A television placed atop the enclosure displayed the video feed from a “lipstick camera” mounted inside the enclosure. The camera had a similar line of sight as the participant, and the participant performed the task while watching the TV.



Figure 3 – Real Space Environment, the user conducted the task while watching the TV.

Purely Virtual Environment (PVE). Participants stood at a four-foot high table, and wore a Virtual Research V8 head-mounted display (HMD) with HiBall tracker and Fakespace Pinchgloves, tracked with Polhemus Fastrak trackers (Figure 4). The self-avatar’s appearance was generic (neutral gray).

The participant picked up a virtual block by pinching two fingers together (i.e. thumb and index finger). The block closest to an avatar’s hand was highlighted to inform the participant which block would be selected by pinching. Pinching caused the virtual block to snap into the virtual avatar’s hand, and the hand appeared to be holding the block. To rotate the block, the participant rotated his hand while maintaining the pinching gesture. When the participant opened the pinch, the virtual block was released, and an open hand avatar was displayed.



Figure 4 – Purely Virtual Environment, virtual objects with a generic self-avatar.

If the block was released within six inches of the workspace surface and over the target grid, it snapped into an unoccupied position in grid. Releasing the block away from the grid caused it to simply drop onto the table. Releasing the block more than six inches above the table caused the block to float in mid-air to aid in rotation. While not realistic, these behaviors, reduced the need for fine-grained interaction, which would have artificially inflated the time to complete the task. There was no inter-block collision detection, and block interpenetration was not automatically resolved.



Figure 5 – Hybrid Environment, real objects with a generic self-avatar.

Hybrid Environment (HE). Participants wore identical yellow dishwashing gloves and the HMD (Figure 5). Within the VE, participants handled physical blocks, identical to the RSE blocks, and saw a self-avatar with accurate shape and generic appearance.

Visually-Faithful Hybrid Environment (VFHE). Participants wore only the HMD. Otherwise, this condition similar to the HE. The self-avatar was visually faithful, as the shape reconstruction was texture-mapped with images from a HMD mounted camera. The participant saw an image of his own hands (Figure 6).




Figure 6 – Visually Faithful Hybrid Environment, real objects and a visually faithful self-avatar.

Virtual Environment. The VE room was identical in all three of the virtual conditions (PVE, HE, VFHE). It had several virtual objects, including a lamp, plant, and painting, along with a virtual table that was registered with a real Styrofoam table. The enclosure in the RSE was rendered with transparency in the VE (Figure 7).



Figure 7 – The VE in all three virtual conditions.

The participant wore a V8 HMD (640x480 resolution in each eye) that was tracked with the original laboratory version of the UNC HiBall tracking system. All the VE conditions were rendered on an SGI Reality Monster. The PVE ran on one rendering pipe at twenty FPS. The HE and VFHE ran on four rendering pipes at twenty FPS for virtual objects, and twelve FPS for reconstructing real

objects. The reconstruction system used 4 NTSC resolution cameras and generated a 320x240 reconstruction. We estimated latency at 0.3 seconds and with 1 cm reconstruction error.

We limited VE time to fifteen minutes, as many pilot subjects complained of fatigue after that amount of time.

Rationale for Conditions. We expected a participant's RSE performance to be the fastest, as the interaction and visually fidelity were optimal. Thus, we compared how closely a participant's VE task performance was to their RSE task performance.

The RSE was used for task training and as a baseline. The block design task had a learning curve, and doing the task in the RSE allowed participants to become proficient without spending additional time in the VE. Pilot testing showed that four practice patterns were needed to stabilize a participant's task performance.

In a pilot study ($n=20$), participants performed the RSE task on a table without the enclosure, and there was no difference in task performance. The enclosure and camera allowed the RSE to have the reduced field of view and working volume as the VE conditions. Further, the enclosure also allowed us to see if any issues arose from performing an out of direct line-of-sight task. Participants did not report, or appear to exhibit, any issues.

The PVE was a *plausible VE approach* to the block task. As in current VEs, most of the objects were virtual, and interactions were done with specialized equipment and gestures. The difference in task performance between the RSE and PVE corresponded primarily to the impedance of interacting with virtual objects.

The HE evaluated the effect of real objects on task performance. We assumed any interaction hindrances caused by the gloves were minor compared to the effect of handling real objects.

The VFHE evaluated the cumulative effect of both real object interaction and visually faithful self-avatars on performance.

3.3. Measures

Task Performance. Participants were timed on replicating *correctly* the target pattern. We also recorded if the participant incorrectly concluded that target pattern was replicated. In these cases, the participant was informed of the error and continued to work. Each participant eventually completed every pattern correctly.

Other Factors. We measured sense-of-presence, spatial ability and simulator sickness by using the Steed-Usoh-Slater Presence Questionnaire (SUS) [12], Guilford-Zimmerman Aptitude Survey Part 5: Spatial Orientation, and the Kennedy – Lane Simulator Sickness Questionnaire respectively. We performed post-

experiment interviews with each participant to gather further data on their impressions and recorded self-reported and experimenter-reported behaviors.

3.4. Experimental Procedure

All participants first completed a consent form and questionnaires to gauge their physical and mental condition, simulator sickness, and spatial ability.

Real Space. Next, the participant entered the room with the real space environment (RSE) setup. The participant was presented with the wooden blocks and was instructed on the task. The participant was also told that they would be timed, and to examine the blocks and become comfortable with moving them. The cloth on the enclosure was lowered, and the TV turned on.

The participant was given a series of six practice patterns, three small (2x2) and then three large (3x3). The participant was told the number of blocks involved in a pattern, and to notify the experimenter when they were done. After the practice patterns were completed, a series of six test patterns (three small, three large) were administered. Between patterns, the participant was asked to randomize the blocks' orientations.

We recorded the time required to complete each test pattern correctly. If the participant misjudged the completion of the pattern, we noted this as an error and told the participant that the pattern was not yet complete, and to continue working on the pattern. The clock was not stopped on errors. The total time was used as the task performance measure for a pattern.

Virtual Space. Next, the participant entered a different room where the experimenter helped the participant put on the HMD and any additional equipment particular to the VE condition (PVE – tracked pinch gloves, HE – dishwashing gloves).

Following a period of adaptation to the VE, the participant practiced on two small and two large patterns. The participant then was timed on two small and two large test patterns. A participant could ask questions and take breaks between patterns if so desired. Only one person (a PVE participant) asked for a break. All finished their VE condition within the fifteen minute time limit.

Post Experience. Finally, the participant was interviewed about their impressions of and reactions to the session. The debriefing session was a semi-structured interview. The questions asked were starting points, and the interviewer could delve into responses for further clarification or to explore unexpected comments.

The participant filled out the simulator sickness questionnaire again. By comparing their pre- and post-experience scores, we could assess if their level of

simulator sickness had changed while performing the task. Finally, an expanded Slater – Usoh – Steed Virtual Presence Questionnaire was given to measure the participant’s sense of presence in the VE.

Managing Anomalies. If the head or hand tracker lost tracking or crashed, we quickly restarted the system (about 5 seconds). We noted long or repeated tracking failures. In almost all the cases, the participants were so engrossed with the task they never *noticed* any problems and continued working. Tall participants were allowed to sit to perform the task to improve head tracker reliability. None of the tracking failures appeared to significantly affect the task performance time.

There were two additional patterns for replacement of voided trials, such as if a participant dropped a block onto the floor. This happened twice and was noted.

3.5. Hypotheses

Participants who manipulate real objects in the VE (HE, VFHE) will complete the spatial cognitive manual task significantly closer to their RSE task performance time than will participants who manipulate virtual objects (PVE). *Handling real objects in VEs improves task performance.*

Participants represented in the VE by a visually faithful self-avatar (VFHE) will complete the spatial cognitive manual task significantly closer to their RSE task performance time than will participants who are represented by a generic self-avatar (PVE, HE). *Self-avatar visual fidelity improves task performance in VEs.*

4. Results

Forty participants completed the study. There were thirteen in both the purely virtual environment (PVE) and hybrid environment (HE) conditions, and fourteen in the visually-faithful hybrid environment (VFHE) condition. The participants were primarily male (thirty-three) undergraduate students enrolled at UNC-CH (thirty-one). Participants were recruited from UNC-CH Computer Science classes and word of mouth.

On a [1..7] scale, participants reported little prior VE experience (M=1.37, s.d.=0.66), high computer usage (M=6.39, s.d.=1.14), and moderate – 1 to 5 hours a week – computer/video game play (M=2.85, s.d.=1.26). There were no significant differences on these measures between the groups.

During the recruiting process, we required participants to have taken or be currently enrolled in a higher-level mathematics course (equivalent of a Calculus 1 course). This reduced spatial ability variability, and in turn reduced task performance variability.

4.1. Experiment Data

The dependent variable for task performance was the difference in the time to correctly replicate target patterns in the VE condition and the time replicate patterns in the RSE. We used a two-tailed t-test with unequal variances and $\alpha=0.05$ level for significance unless otherwise stated.

Figure 8 and Table 1 show the data. Recall that in the RSE, participants worked on three small and three large patterns, and in the VE condition only two of each.




Figure 8 – Time to replicate the target pattern.

Table 1 – Task performance

	Small Pattern Time (seconds)		Large Pattern Time (seconds)	
	Mean	S.D.	Mean	S.D.
RSE (n=40)	16.81	6.34	37.24	8.99
PVE (n=13)	47.24	10.43	116.99	32.25
HE (n=13)	31.68	5.65	86.83	26.80
VFHE (n=14)	28.88	7.64	72.31	16.41

Table 2 – Between groups task performance

	Small Pattern		Large Pattern	
	t-test	p-value	t-test	p-value
PVE-RSE vs. VFHE-RSE	3.32	0.0026**	4.39	0.00016***
PVE-RSE vs. HE-RSE	2.81	0.0094**	2.45	0.021*
VFHE-RSE vs. HE-RSE	1.02	0.32	2.01	0.055+

Significant at the * $\alpha=0.05$, ** $\alpha=0.01$, *** $\alpha=0.001$

+ - requires further investigation

4.2. Other Factors

Sense-of-presence, simulator sickness, and spatial ability were not significantly different between conditions. A full analysis of the sense-of-presence results (Table 3) is beyond the scope of this paper. The correlation between spatial ability and performance was $r = -0.31$ for small patterns, and $r = -0.38$ for large patterns.

Table 3 – Between groups sense of presence, simulator sickness and spatial ability

		PVE vs. VFHE	PVE vs. HE	HE vs. VFHE
Sense-of-Presence	t-test	1.10	1.64	0.64
	p-value	0.28	0.11	0.53
Simulator Sickness	t-test	1.16	0.49	-0.57
	p-value	0.26	0.63	0.58
Spatial Ability	t-test	-1.58	-1.41	0.24
	p-value	0.13	0.17	0.82

5. Discussion

As expected, performing the block-pattern task took longer in any VE than it did in the RSE. The PVE participants took about three **times** as long as they did in the RSE. The HE and VFHE participants took about twice as long as they did in the RSE.

For small and large patterns, both VFHE and HE task performances were significantly better than PVE task performance (Table 2, first two rows). The difference in task performance between the HE and VFHE was not significant at the $\alpha=0.05$ level (Table 2, third row).

We accept the task performance hypothesis; interacting with real objects significantly improved task performance over interacting with virtual objects.

Participants were asked how well they thought they achieved the task, on a scale from 1 (not very well) to 7 (very well). VFHE participants self-reported their performance significantly higher ($M=5.43$, $s.d.=1.09$) than PVE ($M=4.57$, $s.d.=0.94$) participants ($p=0.0345$).

For the case we investigated, *interacting with real objects provided a quite substantial performance improvement over interacting with virtual objects for cognitive manual tasks*. Although task performance in all the VE conditions was substantially worse than in the RSE, the task performance of HE and VFHE participants was significantly better than for PVE participants.

There is a near significant difference between HE and VFHE large pattern performance (Table 2, $p=0.055$), and we do not have a hypothesis as to the cause of this result, particularly when considered with the small pattern result ($p=0.32$). This is a candidate for further investigation.

Although interviews showed visually faithful self-avatars (VFHE) were preferred, there was no statistically significant difference in task performance compared to those presented a generic self-avatar (HE and PVE).

We reject the self-avatar visual-fidelity hypothesis; a visually faithful self-avatar did not improve task performance in a VE, compared to a generic self-avatar.

5.1. Debriefing Trends

Analysis of the post-experience interviews resulted in the following trends:

- Among the participants using the reconstruction system (HE and VFHE), 75% noticed reconstruction errors and 25% noticed the lag. Most complained of the limited field of view. Interestingly, the RSE had a similar field of view, but no participant mentioned it.
- 93% of the PVE and 13% of the HE and VFHE participants felt the interaction was unnatural.
- 25% of the HE and VFHE participants felt the interaction was natural.
- 65% of VFHE and 30% of HE participants commented that their self-avatar “looked real”.
- 43% of PVE participants commented on the blocks not being there or not behaving naturally.

VFHE participants reported feeling comfortable with performing the task significantly more quickly than PVE participants ($T_{26} = 2.83$, $p=0.0044$). VFHE participants were comfortable with the workings of the VE almost an entire practice pattern earlier (1.50 to 2.36 patterns).

5.2. Observations

Time Use: The time to rotate the block dominated the difference in times between VE conditions. The second most significant component of the total time was the selection and placement of the blocks. These factors were improved through the tactile feedback, natural interaction, and motion constraints of handling real blocks. When estimating out the reconstruction lag in the VFHE and HE conditions, performances in those conditions were very close to RSE performance.

Pinch Gloves: Using the one-size-fits-all pinch gloves had fit and hygiene issues in the fourteen-participant PVE group.

- Two participants with large hands had difficulty fitting into the gloves.
- Two participants with small hands had difficulty registering pinching actions, as the sensors were then not positioned appropriately.
- One participant became nauseated and was excused mid-way through the experiment. The pinch gloves became moist with his sweat, and this became a hygiene issue for participants.

Cue Conflict: We saw evidence that the cue conflict of making a real pinching motion and seeing a grasping motion in the (hand) avatar to pick up a block in the PVE affected participants’ actions. This difference caused 25% of the participants to forget the pinching metaphor and try a grasping action (which at times did not register with the pinch gloves). If the experimenter observed this behavior, he made note of it, and reminded the participant to make pinching motions to grasp a block.

Adaptive Behavior: The PVE had some interaction shortcuts for common tasks, e.g. blocks floating or

snapping into place. Some participants, in an effort to maximize efficiency, learned to grab blocks and place them in midair before the beginning of a pattern. This allowed easy and quick access to blocks. The inclusion of the shortcuts was carefully considered to assist in interaction, yet led to adaptation and learned behavior.

Typically participants mentally subdivided the target pattern and matched one subsection at a time. Each block was picked up and rotated until the desired face was found. Some RSE participants noted that this rotation could be done so quickly that they randomly spun each block until they found the desired face. In contrast, two PVE and one HE participant remarked that the slower interaction of block rotation in the VE influenced them to memorize the relative orientation of the block faces to improve performance. For training applications, having participants develop VE-specific behaviors, inconsistent with their real world approach to the task, could be detrimental to effectiveness or even dangerous.

Manipulating real objects also benefited from motion constraints. Tasks such as placing the center block in a nine-block pattern and closing gaps between blocks were easy with real blocks. In the PVE, these interaction tasks would have been difficult and time-consuming. We provided snapping upon release of a block to alleviate these handicaps, but these aides might be questioned if the system was used for learning or training.

6. Conclusions

Interacting with real objects significantly improves task performance over interacting with virtual objects in spatial cognitive tasks. Importantly, it brings performance closer to that of doing the task in real space.

Handling real objects makes task interaction in the VE more like the actual real-world task. Even in our simple task, we saw evidence that manipulating virtual objects sometimes caused participants to develop VE-specific approaches. Training and simulation VEs attempt to recreate real experiences and would benefit from having the participant manipulate as many real objects as possible. If an object reconstruction is not used, we feel instrumenting, modeling and tracking the real objects that the participant will handle, would enhance performance and learning spatial cognitive tasks.

Self-avatar visual fidelity is clearly secondary to interacting with real objects, and probably has little, if any, affect on cognitive task performance. We believe that a visually faithful self-avatar is better than a generic self-avatar, but from a task performance standpoint, the advantages do not seem very strong.

We know that the purely virtual aspect of current VEs has limited the applicability to some tasks. In future work, we hope to identify the tasks that would most benefit from having the participant handle real objects.

7. Acknowledgements

We thank the study participants, Edward Johnson and Carolyn Kanagy (UNC-Chapel Hill Department of Psychology), Henry Fuchs, Herman Towles, David Harrison, Kurtis Keller, John Thomas, and the UNC-Chapel Hill Effective Virtual Environments group. This work was supported by the LINK Simulation and Training Foundation, ONR Virte Program, NIH National Institute of Biomedical Imaging and Bioengineering, NSF Information Technology Research BioGeometry Project, UNC-Chapel Hill, and UNC-Charlotte.

8. Bibliography

- [1] F. Brooks "What's Real About Virtual Reality?" *IEEE Computer Graphics and App.*, Vol 19, No 6, pp 16-27.
- [2] I. Sutherland. "The Ultimate Display", In *Proceedings of IFIP 65*, Vol 2, pp 506, 1965.
- [3] M. Slater and M. Usoh. "The Influence of a Virtual Body on Presence in Immersive Virtual Environments", *VR 93, Virtual Reality International, Proceedings of the Third Annual Conference on Virtual Reality*, London, Meckler, 1993, pp 34-42.
- [4] M. Usoh, K. Arthur, *et al.* "Walking > Virtual Walking> Flying, in Virtual Environments", *Proceedings of SIGGRAPH 99*, pp 359-364, Computer Graphics Annual Conference Series, 1999.
- [5] A. Hilton, D. Beresford, T. Gentils, R. Smith, W. Sun, and J. Illingworth. "Whole-Body Modeling of People from Multiview Images to Populate Virtual Worlds", *The Visual Computer*, Vol 16, No 7, pp 411-436, 2000.
- [6] H. Hoffman, A. Carlin, and S. Weghorst. "Virtual Reality and Tactile Augmentation in the Treatment of Spider Phobia", *Medicine Meets Virtual Reality 5*, San Diego, California, January, 1997.
- [7] K. Hinckley, R. Pausch, J. Goble, and N. Kassell. "Passive Real-World Interface Props for Neurosurgical Visualization", In *Proceedings of the 1994 SIG-CHI Conference*, pp 452-458.
- [8] D. Bowman and L. Hodges. "An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments", M. Cohen and D. Zeltzer, Eds., In *Proceedings 1997 ACM Symposium on Interactive 3-D Graphics*, April 1997, pp 35-38.
- [9] C. Hand. A Survey of 3-D Interaction Techniques, *Computer Graphics Forum*, Vol 16, No 5, pp 269-281 (1997). Blackwell Publishers. ISSN 1067-7055.
- [10] B. Lok. "Online Model Reconstruction for Interactive Virtual Environments", In *Proceedings 2001 ACM Symposium on Interactive 3-D Graphics*, Chapel Hill, N.C., 18-21, March 2001, pp 69-72, 248.
- [11] D. Wechsler. The Measurement of Adult Intelligence. 1st Ed., Baltimore, MD: Waverly Press, Inc. 1939.
- [12] M. Usoh, E. Catena, S. Arman, and M. Slater. "Using Presence Questionnaires in Reality" [*Presence: Teleoperators and Virtual Environments*](#), Vol 9, No 5, pp 497-503.