

# The Effect of 3D Widget Representation and Simulated Surface Constraints on Interaction in Virtual Environments

Robert W. Lindeman<sup>1</sup> John L. Sibert<sup>1</sup> James N. Templeman<sup>2</sup>

<sup>1</sup>*Department of Computer Science  
The George Washington University  
801 22nd Street, NW, Washington, DC, 20052  
+1-202-994-7181  
[gogo | sibert]@seas.gwu.edu*

<sup>2</sup>*Information Technology Division  
Naval Research Laboratory, Code 5510  
Washington, DC, 20375-5337  
+1-202-767-5613  
templeman@itd.nrl.navy.mil*

## Abstract

*This paper reports empirical results from two studies of effective user interaction in immersive virtual environments. The use of 2D interaction techniques in 3D environments has received increased attention recently. We introduce two new concepts to the previous techniques: the use of 3D widget representations; and the imposition of simulated surface constraints. The studies were identical in terms of treatments, but differed in the tasks performed by subjects. In both studies, we compared the use of two-dimensional (2D) versus three-dimensional (3D) interface widget representations, as well as the effect of imposing simulated surface constraints on precise manipulation tasks. The first study entailed a drag-and-drop task, while the second study looked at a slider-bar task. We empirically show that using 3D widget representations can have mixed results on user performance. Furthermore, we show that simulated surface constraints can improve user performance on typical interaction tasks in the absence of a physical manipulation surface. Finally, based on these results, we make some recommendations to aid interface designers in constructing effective interfaces for virtual environments.*

## 1. Introduction

Previous research into the development of usable Virtual Environment (VE) interfaces has shown that applying 2D techniques can be effective for tasks involving widget manipulation [12, 5, 6]. It has also been shown that the use of physical props for manipulating these widgets can provide effective feedback to aid the

user [13, 11]. In some cases, however, it is not practical to require the use of physical props because of real-world constraints, such as a cramped workspace, or the need to use both hands for another task. This paper presents empirical results from experiments designed to identify those aspects of 2D interfaces used in 3D spaces that have the greatest influence on effectiveness for manipulation tasks.

## 2. Previous Work

A number of novel techniques have been employed to support interaction in 3-space. Glove interfaces allow the user to interact with the environment using gestural commands [9, 8, 3]. Laser-pointer techniques provide menus that float in space in front of the user, and are accessed using either the user's finger or a 3D mouse [14, 15, 6]. With these types of interfaces, however, it is difficult to perform precise movements, such as dragging a slider to a specified location, or selecting from a pick list. Part of the difficulty in performing these tasks comes from the fact that the user is pointing in free space, without the aid of anything to steady the hands [14].

Another major issue with the floating windows interfaces comes from the inherent problems of mapping a 2D interface into a 3D world. One of the reasons the mouse is so effective on the desktop, is that it is a 2D input device used to manipulate 2D (or 2.5D) widgets on a 2D display. Once we move these widgets to 3-space, the mouse is less tractable as an input device. Wloka [16] and Deering [7] attempt to solve this problem by using menu widgets that pop up in a fixed location relative to a 3D mouse. With practice, the user learns where the menu is in relation to the mouse, so the depth can be learned. Though

effective for simple menu selection tasks, these methods provide limited user precision for more complex tasks because of a lack of physical support for manipulations.

To counter this, some researchers have introduced the use of "pen-and-tablet" interfaces [1, 2, 14, 10, 12]. These approaches register an interaction window with a physical prop held in the non-dominant hand. The user interacts with these hand-held windows using either a finger, or a stylus held in the dominant hand. This type of interface combines the power of 2D window interfaces with the necessary freedom of movement provided by 3D interfaces. Though the promise of these interfaces has been outlined [14, 13, 5], limited empirical data has been collected to measure the effectiveness of these interfaces [11, 2, 13]. Here we present results that aim to refine our knowledge of the nature of these interaction techniques.

## 2.1. Types of Actions

We have designed a testbed which provides 2D widgets for testing typical UI tasks, such as drag-and-drop, manipulating slider-bars, and pressing buttons, from within a 3D environment. Our system employs a virtual paddle, registered with a physical paddle prop, as a work surface for manipulating 2D interface widgets within the VE. Simple collision detection between the tracked index fingertip of the user and interface widgets is used for selection.


In our previous work [13], we describe two studies comparing unimanual versus bimanual manipulation, and the presence of passive-haptic feedback versus having no haptic feedback. In the current work, we looked at ways of improving non-haptic interfaces for those systems where it is impractical to provide haptic feedback. The two possible modifications we studied are the use of 3D representations of widgets instead of 2D representations, and the imposition of simulated physical surface constraints.

We chose two different types of tasks in order to study representative 2D manipulations. A drag-and-drop type of task required subjects to move a shape from a start to a target position. A one-dimensional slide-bar task required subjects to adjust a slider within some threshold to match a target position.

## 2.2. Widget Representation

In our previous studies [13], all the shapes that were manipulated were 2D, flush with the plane of the work surface. Even though the shapes had a 3D bounding volume for detecting collisions, only a 2D shape was displayed to the user. One could argue that this was optimized more for the treatments where a physical surface was present than where no physical surface was

present. In our experiments, we compared the use of these 2D representations with shapes that had depth, providing additional visual feedback as to the extent of the bounding volume (Figure 1). The idea is that by providing subjects with visual feedback as to how deep the fingertip was penetrating the shape, they would be able to maintain a constant depth more easily, improving performance [4]. This would allow statements to be made about the influence of visual widget representation on performance and preference measures.



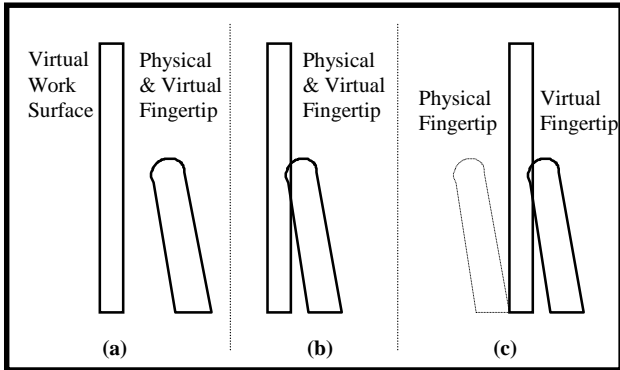
**Figure 1:** 3D widget representation. Shape is solid; Target is wireframe

## 2.3. Simulated Surface Constraints

The superiority of the passive-haptic treatments over the non-haptic treatments from the experiments presented in [13] led us to analyze which aspects of the physical surface accounted for its superiority. The presence of a physical surface: 1) provides tactile feedback felt by the dominant-hand index fingertip; 2) provides haptic feedback felt in the extremities of the user, steadying movements in a way similar to moving the mouse resting on a tabletop; and 3) constrains the motion of the finger along the Z axis of the work surface to lie in the plane of the surface, thereby making it easier for users to maintain the necessary depth for selecting shapes.

In order to differentiate between the amount each of these aspects influences overall performance, the notion of *clamping* is introduced. Clamping involves imposing a simulated surface constraint to interfaces that do not provide a physical work surface (Figure 2). During interaction, when the real finger passes a point where a physical surface would be (if there were a physical surface), the virtual finger avatar is constrained such that the fingertip remains intersected with the work surface

avatar. Movement in the X/Y-plane of the work surface is unconstrained; only the depth of the virtual fingertip is constrained. If the subject presses the physical finger past a threshold depth, the virtual hand pops through the surface, and is registered again with the physical hand.



**Figure 2:** Clamping (a) Fingertip approaches work surface; (b) Fingertip intersects work surface; (c) Virtual fingertip clamped to work surface

We hypothesized that clamping would make it easier for subjects to keep the shapes selected during manipulation, even if no haptic feedback was present, because it would be easier to maintain the necessary depth. It should make object grasping easier, and make the interaction more like it is in real life.

### 3. Experimental Method

This section describes the experimental design used in our experiments. These experiments were designed to compare interfaces that combine *2D* or *3D* widget representations with the three surface types of *physical*, *clamped*, and *none*. The case where no physical surface is present is interesting, in that some systems employ purely virtual tools, possibly chosen from a tool bar, using a single, common handle. By virtualizing the head of the tool, the functionality of a single handle can be overloaded, providing greater flexibility (with the possible trade-off of increased cognitive load).

We use quantitative measures of proficiency, such as mean task completion time and mean accuracy, as well as qualitative measures, such as user preference, to compare the interfaces. Two experiments, one involving a drag-and-drop task, and one involving the manipulation of a slider-bar, were conducted. In the interest of brevity, we present them together.

#### 3.1. Experimental Design

The experiments were designed using a  $2 \times 3$  factorial within-subjects approach, with each axis representing one

independent variable. The first independent variable was the use of *2D* widget representations (2) or *3D* widget representations (3). The second independent variable was surface type: physical (P); clamped (C); or no surface was present (N).

Six different interaction techniques (treatments) were implemented which combine these two independent variables into a  $2 \times 3$  matrix, as shown in Table 1.

	<b>2D Widget Representation (2)</b>	<b>3D Widget Representation (3)</b>
<b>Physical Surface (P)</b>	2P Treatment	3P Treatment
<b>Clamped Surface (C)</b>	2C Treatment	3C Treatment
<b>No Surface (N)</b>	2N Treatment	3N Treatment

**Table 1:**  $2 \times 3$  design

Each cell is defined as:


- 2P** = 2D Widget Representation, Physical Surface
- 2C** = 2D Widget Representation, Clamped Surface
- 2N** = 2D Widget Representation, No Surface
- 3P** = 3D Widget Representation, Physical Surface
- 3C** = 3D Widget Representation, Clamped Surface
- 3N** = 3D Widget Representation, No Surface

For the **2P** treatment, subjects held a paddle-like object in the non-dominant hand (Figure 3), with the work surface defined to be the face of the paddle. Subjects could hold the paddle in any position that felt comfortable and allowed them to accomplish the tasks quickly and accurately. Subjects were presented with a visual avatar of the paddle that matched exactly the physical paddle in dimension (Figure 4). For the **2C** treatment, the subjects held only the handle of the paddle in the non-dominant hand (no physical paddle head), while being presented with a full paddle avatar. The same visual feedback was presented as in 2P, but a clamping region was defined just behind the surface of the paddle face. The clamping region was a box with the same X/Y dimensions as the paddle surface, and a depth of 3cm (determined from a pilot study). When the real index fingertip entered the clamp region, the hand avatar was "snapped" so that the virtual fingertip was on the surface of the paddle avatar. For the **2N** treatment, the user held the paddle handle (no physical paddle head) in the non-dominant hand, but was presented with a full paddle avatar in the VE. The only difference between 2C and 2N was the lack of clamping in 2N. The avatar hand could freely pass into and through the virtual paddle.



**Figure 3:** The physical paddle

The **3P**, **3C**, and **3N** treatments were identical to **2P**, **2C**, and **2N**, respectively, except for the presence of 3D widget representations. The widgets were drawn as volumes, as opposed to polygons, with the back side of the volume flush with the paddle surface, and the front side extending forward 0.8cm (determined from a pilot study). The widgets were considered selected when the fingertip of the hand avatar intersected the bounds of the volume.




**Figure 4:** The docking task

Using a diagram-balanced Latin squares approach to minimize ordering effects, six different orderings of the treatments were defined, and subjects were assigned at random to one of the six orderings. Subjects were randomly assigned to one of two groups. Each group performed 20 trials on one of two tasks over each treatment. Six different random orderings for the 20 trials were used. The subjects were seated during the entire experiment.

Task one was a docking task. Subjects were presented with a colored shape on the work surface, and had to slide it to a black outline of the same shape in a different location on the work surface, and then release it (Figure 4). Subjects could repeatedly adjust the location of the shape until they were satisfied with its proximity to the outline shape, and then move on to the next trial by pressing a "Continue" button displayed in the center at the lower edge of the work surface. This task was designed to test the component UI action of "Drag-and-Drop," which is a continuous task. The trials were a mix between horizontal, vertical, and diagonal movements.

The second task was a one-dimensional sliding task. The surface of the paddle displayed a slider-bar and a number (Figure 5). The value of the number, which could range between 0 and 50, was controlled by the position of the slider-pip. A signpost was displayed in the VE, upon which a target number between 0 and 50 was displayed. The signpost was positioned directly in front of the subject. The subject had to select the slider-pip, and slide it along the length of the slider, until the number on the paddle matched the target number on the signpost, release the slider-pip, and then press the "Continue" button to move on to the next trial. The subject could adjust the slider-pip before moving on to the next trial. This task was designed to test the component UI action of "Slider-bar," which is also a continuous task.




**Figure 5:** The sliding task

The trials were a mix between horizontal and vertical sliders (Figure 6). The slider-bar was 14cm long and 3cm thick for both the horizontal and vertical trials. This gave a slider sensitivity for user movement of 0.3cm per number (determined from a pilot study). This means the pip had a tolerance of 0.3cm before the number on the paddle would change to the next number.

### 3.2. Shape Manipulation

For the sliding task, the slider-pip was red, the number on the paddle was yellow, and the number on the signpost

was green. Subjects selected shapes (or the pip) simply by moving the fingertip of their dominant-hand index finger to intersect it. The subject would release by moving the finger away from the shape, so that the fingertip no longer intersected it. For the docking task, this required the subject to lift (or push) the fingertip so that it no longer intersected the virtual work surface, as moving the finger tip along the plane of the work surface translated the shape along with the fingertip. This was true for the sliding task as well, except that sliding the finger away from the long axis of the slider also released the pip. Once the fingertip left the bounding box of the shape (or pip), it was considered released.



**Figure 6:** Paddle layout for the sliding task; (a) Horizontal bar; (b) Vertical bar (dashed lines are widget positions for left-handed subjects)

### 3.3. System Characteristics

Our software was running on a two-processor SGI Onyx workstation equipped with a RealityEngine<sup>2</sup> graphics subsystem, two 75MHz MIPS R8000 processors, 64 megabytes of RAM, and 4 megabytes of texture RAM. Because of a lack of audio support on the Onyx, audio feedback software (see below) was run on an SGI Indy workstation, and communicated with our software over Ethernet. The video came from the Onyx, while the audio came from the Indy. We used a Virtual I/O i-glasses HMD in monoscopic mode to display the video and audio. The positions and orientations of the head, index finger, and paddle were gathered using 6-DOF trackers. The software used the OpenInventor library from SGI, and ran in one Unix thread. A minimum of 11 frames per second (FPS) and a maximum of 16 FPS were maintained throughout the tests, with the average being 14 FPS.

### 3.4. Subject Demographics

Thirty-six subjects for each task (72 total) were selected on a first-come, first-served basis, in response to a call for subjects. For the docking task, most of the subjects were college students (22), either undergraduate (10) or graduate (12). The rest (14) were not students. The

mean age of the subjects was 30 years, 5 months. In all, 33 of the subjects reported they used a computer at least 10 hours per week, with 25 reporting computer usage exceeding 30 hours per week. The remaining 3 subjects reported computer usage between 5 and 10 hours per week. Five subjects reported that they used their left hand for writing. Thirteen of the subjects were female and 23 were male. Fifteen subjects said they had experienced some kind of "Virtual Reality" before. All subjects passed a test for colorblindness. Nine subjects reported having suffered from motion sickness at sometime in their lives, when asked prior to the experiment.

For the sliding task, most of the subjects were college students (26), either undergraduate (15) or graduate (11). The rest (10) were not students. The mean age of the subjects was 27 years, 5 months. In all, 30 of the subjects reported they used a computer at least 10 hours per week, with 21 reporting computer usage exceeding 30 hours per week. Of the remaining 6 subjects, 3 reported computer usage between 5 and 10 hours per week, and 3 between 1 and 5 hours per week. Four subjects reported that they used their left hand for writing. Thirteen of the subjects were female and 23 were male. Twenty-four subjects said they had experienced some kind of "Virtual Reality" before. All subjects passed a test for colorblindness. Ten subjects reported having suffered from motion sickness at sometime in their lives, when asked prior to the experiment.

### 3.5. Protocol

Before beginning the actual experiment, the demographic information reported above was collected. The user was then fitted with the dominant-hand index finger tracker, and asked to adjust it so that it fit snugly. Each subject was read a general introduction to the experiment, explaining what the user would see in the virtual environment, which techniques they could use to manipulate the shapes in the environment, how the paddle and dominant-hand avatars mimicked the motions of the subject's hands, and how the HMD worked.

After fitting the subject with the HMD, the software was started, the visuals appeared, and the audio emitted two sounds. The subjects were asked if they heard the sounds at the start of each task. To help subjects orient themselves, they were asked to look at certain virtual objects placed in specific locations within the VE, and this process was repeated before each treatment. The user was then given the opportunity to become familiar with the dominant-hand and paddle avatars.

The work surface displayed the message, 'To begin the first trial, press the "Begin" button.' Subjects were asked to press the "Begin" button on the work surface by touching it with their finger. After doing this, they were

given practice trials, during which a description of the task they had to perform within the VE was given. The user was given as many practice trials as they wanted, and instructed that after practicing, they would be given 20 more trials which would be scored in terms of both time and accuracy. They were instructed to indicate when they felt they could perform the task quickly and accurately given the interface they had to use. The subject was coached as to how best to manipulate the shapes, and about the different types of feedback they were being given. For instance, for the clamping treatments (2C & 3C), a detailed description of what clamping is was given.

After the practice trials, the subject was asked to take a brief rest, and was told that when ready, 20 more trials would be given, and would be scored in terms of both time and accuracy. It was made clear to the subjects that neither time nor accuracy was more important than the other, and that they should try to strike a balance between the two. Accuracy for the docking task was measured by how close the center of the shape was placed to the center of the target position, and for the sliding task, accuracy was measured as how closely the number on the paddle matched the target number on the signpost. After each treatment, the HMD was removed, the paddle was taken away, and the subject was allowed to relax as long as they wanted to before beginning the next treatment.

### 3.6. Additional Feedback

In addition to visual and (in some cases) haptic feedback, our system provided other cues for the subject, regardless of treatment. First, the tip of the index finger of the dominant-hand avatar was colored yellow, to give contrast against the paddle surface. Second, in order to simulate a shadow of the dominant hand, a red drop-cursor, which followed the movement of the fingertip in relation to the plane of the paddle surface, was displayed on the work surface. When the fingertip was not in the space directly in front of the work surface, no cursor was displayed. To help the subjects gauge when the fingertip was intersecting UI widgets, each widget became highlighted, and an audible *CLICK!* sound was output to the headphones worn by the subject. When the user released the widget, it returned to its normal color, and a different *UNCLICK!* sound was triggered. For the sliding task, each time the number on the paddle changed, a sound was triggered to provide extra feedback to the subject.

Several issues arose for the clamping treatments during informal testing of the technique. One of the problems with the use of clamping is the discontinuity in the mapping of physical-to-virtual finger movement it introduces into the system. This manifests itself in several ways in terms of user interaction. First, because during

clamping the physical and virtual fingertips are no longer registered, lifting the finger from the surface of the paddle (a movement in the Z direction) does not necessarily produce a corresponding movement in the virtual world, as long as the movement occurs solely within the clamping area. This makes releasing the shapes difficult (the opposite problem of what clamping was designed to solve!). This issue was addressed by introducing prolonged practice and coaching sessions before each treatment.

A second problem is the inability of users to judge how "deep" their physical fingertip is through the surface. Even if subjects understand the movement mapping discontinuity, judging depth can still be a problem. To counter this, the fingertip of the index finger, normally yellow, was made to change color, moving from orange to red, as a function of how deep the physical finger was past the point where a physical surface would be if there were one. Again, substantial practice and coaching was given to allow subjects to master this concept. To summarize, clamping consisted of constraining the virtual fingertip to lie on the surface of the paddle avatar, and varying the fingertip color as a function of physical fingertip depth past the (non-existent) physical paddle surface.

### 3.7. Data Collection

Qualitative data was collected for each treatment using a questionnaire. Five questions, arranged on Likert scales, were administered to gather data on perceived ease-of-use, appeal, arm fatigue, eye fatigue, and motion sickness, respectively. The questionnaire was administered after each treatment.

Quantitative data was collected by the software for each trial of each task. The type of data collected was similar for the two tasks. For the docking task, the start position, target position, and final position of the shapes were recorded. For the sliding task, the starting number, target number, and final number were recorded. For both tasks, the total trial time and the number of times the subject selected and released the shape (or pip) for each trial was recorded.

A Fitts-type analysis was not performed for this study. The justification for this comes from the fact that two hands are involved in the interaction techniques which is not represented in Fitts' law.

### 3.8. Results

In order to produce an overall measure of subject preference for the six treatments, we have computed a composite value from the qualitative data. This measure is computed by averaging each of the Likert values from the five questions posed after each treatment. Because

"positive" responses for the five characteristics were given higher numbers, on a scale between one and five, the average of the ease-of-use, appeal, arm fatigue, eye fatigue, and motion sickness questions gives us an overall measure of preference. A score of 1 signifies a lower preference than a score of 5.

The results of the univariate  $2 \times 3$  factorial ANOVA of the performance measures and treatment questionnaire responses for the docking task are shown in Table 2, and those for the sliding task in Table 3. Each row in the tables represents a separate measure, and the mean, standard deviation, f-value, and significance is given for each independent variable. If no significance is found across a given level of an independent variable, then a line is drawn beneath the levels that are statistically equal. The f-value for interaction effects is given in a separate column.

Measure	Widget Representation	Surface Type	Interaction
<b>Docking Time (s)</b>	<b>2D 3D</b>	<b>P C N</b>	
Mean	7.15 6.69	5.32 7.34 8.10	
Stand. Dev.	(2.35) (2.35)	(2.15) (2.44) (2.65)	
Pairwise		P C*** C N** P N***	
f-Value	f = 4.36*	f = 64.34***	f = 0.46
<b>End Distance (cm)</b>	<b>2D 3D</b>	<b>P C N</b>	
Mean	0.13 0.13	0.11 0.13 0.14	
Stand. Dev.	(0.6) (0.05)	(0.06) (0.05) (0.06)	
Pairwise		P C*** C N*** P N***	
f-Value	f = 0.49	f = 7.31***	f = 0.53
<b>Composite Value</b>	<b>2D 3D</b>	<b>P C N</b>	
Mean	3.63 3.66	4.11 3.55 3.27	
Stand. Dev.	(0.52) (0.51)	(0.50) (0.59) (0.61)	
Pairwise		P C*** C N*** P N***	
f-Value	f = 0.75	f = 59.93***	f = 2.42
	df = 1/35	df = 2/70	df = 2/70
*p < 0.05			
**p < 0.01			
***p < 0.001			

**Table 2:**  $2 \times 3$  Factorial ANOVA of performance and subjective measures for the docking task

### 3.9. Discussion

For the docking task, subjects performed faster using 3D widget representations (Docking Time = 6% faster) than with 2D widget representations. Also, subjects performed faster when a physical surface was present (Docking Time = 28% faster) than with clamping, and faster with clamping (Docking Time = 9% faster) than with no surface. There was no difference in accuracy between 3D and 2D widget representations, but accuracy was 15% better with a physical surface than with clamping, and accuracy with clamping was 7% better than with no surface. Looking at the subjective measures, the Composite Preference Value for the main effects shows that subjects had no preference when it came to widget representation, but preferred the physical surface over the clamped surface by 14%, and the clamped surface over no surface by 8%.

For the slider bar task, subjects performed faster using 2D widget representations (Sliding Time = 5% faster)

than 3D widget representations. Also, subjects performed faster when a physical surface was present (Sliding Time = 22% faster) than with clamping, but there was no difference between clamping and no surface. Accuracy was 19% better using 3D widget representations compared to 2D, but there was no difference in accuracy between the physical, clamping, and no surface treatments. Looking at the subjective measures, the Composite Preference Value for the main effects shows that subjects had no preference when it came to widget representation, but preferred the physical surface over the clamped surface by 12%, and the clamped surface over no surface by 5%.

Measure	Widget Representation	Surface Type	Interaction
<b>Sliding Time (s)</b>	<b>2D 3D</b>	<b>P C N</b>	
Mean	5.88 6.21	5.01 6.41 6.72	
Stand. Dev.	(1.36) (1.58)	(1.08) (1.76) (1.67)	
Pairwise		P C*** C N P N***	
f-Value	f = 6.39*	f = 59.77***	f = 1.25
<b>End Distance</b>	<b>2D 3D</b>	<b>P C N</b>	
Mean	0.36 0.29	0.24 0.34 0.38	
Stand. Dev.	(0.29) (0.28)	(0.26) (0.32) (0.41)	
Pairwise		P C C N P N	
f-Value	f = 5.06*	f = 2.44	f = 0.12
<b>Composite Value</b>	<b>2D 3D</b>	<b>P C N</b>	
Mean	3.82 3.88	4.23 3.74 3.57	
Stand. Dev.	(0.45) (0.47)	(0.41) (0.42) (0.60)	
Pairwise		P C*** C N* P N***	
f-Value	f = 0.79	f = 46.22***	f = 0.33
	df = 1/35	df = 2/70	df = 2/70
*p < 0.05			
**p < 0.01			
***p < 0.001			

**Table 3:**  $2 \times 3$  Factorial ANOVA of performance and subjective measures for the sliding task

The addition of a physical surface to VE interfaces can significantly decrease the time necessary to perform UI tasks while increasing accuracy. For those environments where using a physical surface may not be appropriate, we can simulate the presence of a physical surface by using clamping. Our results show some support, though mixed, for the use clamping to improve user performance over not providing such surface intersection cues.

We found that some learning effects were present. On average, subjects were 20% faster on the docking task by the sixth treatment they were exposed to as compared to the first. For the sliding task, a less-dramatic learning effect was present, and seemed to completely disappear by the third treatment. Neither accuracy nor preference values showed learning effects on either task.

There are some subtle differences in the tasks that may account for the differences we saw in the results. First, there was a clear "right answer" for the sliding task (i.e. make the number on the paddle match the target number), whereas it was much more difficult on the docking task to know when the shape was *exactly* lined up with the target. Secondly, the sliding task required the subjects to look around in the VE in order to acquire the target number

from the signpost, whereas with the docking task, the subject only needed to look at the surface of the paddle. Finally, because the sliding task required only one-dimensional movement, along the long axis of the slider, the task was less complex than the docking task, which required subjects to control two degrees of freedom.

#### 4. Observations

The quantitative results from these experiments were clearly mixed. In lieu of drawing conclusions from the work, we convey instead some observations made while conducting them, and offer advice to UI designers of VR systems.

As shown here and in our previous work, using a physical work surface can significantly improve performance. Physical prop use is being reported more frequently in recent literature, and we encourage this trend.

We included complementary feedback in our system (e.g. audio, color variation). Though this might have skewed our results compared to studying a single feedback cue in isolation, we felt that providing the extra feedback, and holding it constant across all treatments, had greater benefit and application than leaving it out. Because life is a multimedia experience, we encourage this approach.

The sensing and delivery technology used in a given system clearly influence how well people can perform using the system. Designers should limit the degree of precision required in a system to that supported by the technology, and apply necessary support (such as snap-grids) where appropriate.

Most VR systems use approaches similar to the 2N or 3N treatments reported here. Also, these systems typically restrict menu interaction to ballistic button presses. We feel that the latter is a consequence of the former. If greater support were given to menu interaction, we might see more complex interactions being employed, allowing greater freedom of expression.

Our future experiments will look at more UI tasks, such as how cascading menus can be effectively accessed from within a VE.

#### 5. Acknowledgements

This research was supported in part by the Office of Naval Research.

#### 6. References

- [1] Billinghamurst, M., Baldis, S., Matheson, L., Philips, M., "3D Palette: A Virtual Reality Content Creation Tool," *Proc. of VRST '97*, (1997), pp. 155-156.
- [2] Bowman, D., Wineman, J., Hodges, L., "Exploratory Design of Animal Habitats within an Immersive Virtual Environment," GA Institute of Technology GVU Technical Report GIT-GVU-98-06, (1998).
- [3] Bryson, S., Levit, C., "The Virtual Windtunnel: An Environment for the Exploration of Three-Dimensional Unsteady Flows," *Proc. of Visualization '91*, (1991), pp. 17-24.
- [4] Conner, B., Snibbe, S., Herndon, K., Robbins, D., Zelesnik, R., van Dam, A., "Three-Dimensional Widgets," *Computer Graphics, Proceedings of the 1992 Symposium on Interactive 3D Graphics*, 25, 2, (1992), pp. 183-188.
- [5] Coquillart, S., Wesche, G., "The Virtual Palette and the Virtual Remote Control Panel: A Device and an Interaction Paradigm for the Responsive Workbench," *Proc. of IEEE Virtual Reality '99*, (1999), pp. 213-216.
- [6] Cutler, L., Fröhlich, B., Hanrahan, P., "Two-Handed Direct Manipulation on the Responsive Workbench," *1997 Symp. on Interactive 3D Graphics*, Providence, RI, (1997), pp. 107-114.
- [7] Deering, M., "The HoloSketch VR Sketching System," *Comm. of the ACM*, 39, 5, (1996), pp. 55-61.
- [8] Fels, S., Hinton, G., "Glove-TalkII: An Adaptive Gesture-to-Format Interface," *Proc. of SIGCHI '95*, (1995), pp. 456-463.
- [9] Fisher, S., McGreevy, M., Humphries, J., Robinett, W., "Virtual Environment Display System," *1986 Workshop on Interactive 3D Graphics*, Chapel Hill, NC, (1986), pp. 77-87.
- [10] Fuhrmann, A., Löffelmann, H., Schmalstieg, D., Gervautz, M., "Collaborative Visualization in Augmented Reality," *IEEE Computer Graphics and Applications*, 18, 4, (1998), pp. 54-59.
- [11] Hinckley, K., Pausch, R., Proffitt, D., Patten, J., Kassell, N., "Cooperative Bimanual Action," *Proc. of SIGCHI '97*, (1997), pp. 27-34.
- [12] Lindeman, R., Sibert, J., Hahn, J., "Hand-Held Windows: Towards Effective 2D Interaction in Immersive Virtual Environments," *Proc. of IEEE Virtual Reality '99*, (1999), pp. 205-212.
- [13] Lindeman, R., Sibert, J., Hahn, J., "Towards Usable VR: An Empirical Study of User Interfaces for Immersive Virtual Environments," *Proc. of the SIGCHI '99*, (1999), pp. 64-71.
- [14] Mine, M., Brooks, F., Séquin, C., "Moving Objects in Space: Exploiting Proprioception in Virtual-Environment Interaction," *Proc. of SIGGRAPH '97*, (1997), pp. 19-26.
- [15] van Teylingen, R., Ribarsky, W., and van der Mast, C., "Virtual Data Visualizer," *IEEE Transactions on Visualization and Computer Graphics*, 3, 1, (1997), pp. 65-74.
- [16] Wloka, M., Greenfield, E., "The Virtual Tricorder: A Uniform Interface for Virtual Reality," *Proc. of UIST '95*, (1995), pp. 39-40.