The Partial-Occlusion Effect: Utilizing Semitransparency in 3D Human-Computer Interaction

SHUMIN ZHAI, WILLIAM BUXTON, and PAUL MILGRAM
University of Toronto

This study investigates human performance when using semitransparent tools in interactive 3D computer graphics environments. The article briefly reviews techniques for presenting depth information and examples of applying semitransparency in computer interface design. We hypothesize that when the user moves a semitransparent surface in a 3D environment, the *partial-occlusion* effect introduced through semitransparency acts as an effective cue in target localization—an essential component in many 3D interaction tasks. This hypothesis was tested in an experiment in which subjects were asked to capture dynamic targets (virtual fish) with two versions of a 3D box cursor, one with and one without semitransparent surfaces. Results showed that the partial-occlusion effect through semitransparency significantly improved users' performance in terms of trial completion time, error rate, and error magnitude in both monoscopic and stereoscopic displays. Subjective evaluations supported the conclusions drawn from performance measures. The experimental results and their implications are discussed, with emphasis on the relative, discrete nature of the partial-occlusion effect and on interactions between different depth cues. The article concludes with proposals of a few future research issues and applications of semitransparency in human-computer interaction.

Categories and Subject Descriptors: H.1.2 [**Models and Principles**]: User/Machine Systems—*human factors*; H.5.2 [**Information Interface and Presentation**]: User Interfaces—*input devices and strategies; interaction styles*; I.3.6 [**Computer Graphics**]: Methodology and Techniques—*interaction techniques*; I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*virtual reality*

General Terms: Design, Experimentation, Human Factors, Measurement

Additional Key Words and Phrases: 3D interfaces, depth perception, partial occlusion, semitransparency, stereopsis

---

Primary support for this work has come from the Information Technology Research Centre of Ontario, the Defense and Civil Institute of Environmental Medicine, the Natural Sciences and Engineering Research Council of Canada, and Xerox Parc, all of whom we gratefully acknowledge.

Authors' addresses: S. Zhai, IBM Almaden Research Center, 650 Harry Road K54/B2, San Jose, CA 95120-6099; email: zhai@almaden.ibm.com; B. Buxton, CSRI, University of Toronto, Toronto, Ontario, Canada, M5S 3H5; email: buxton@aw.sgi.com; P. Milgram, Department of Industrial Engineering, University of Toronto, Toronto, Ontario, Canada M5S 1A4; email: milgram@ie.toronto.edu.

1. INTRODUCTION

With the advent of modern workstations and increasingly high performance personal computers, requirements of efficient manipulation of 3D data are moving from restricted domains to mainstream applications. Examples of such applications include information visualization [Card et al. 1991], virtual environments and telepresence [Ellis et al. 1991; Zeltzer 1992], computer-aided design, telerobotics [Sheridan 1992], and entertainment. With this move to 3D, however, we see a breakdown in many of the interaction techniques that have traditionally been used in 2D direct-manipulation systems. Tasks such as target acquisition, positioning, dragging, sweeping out regions, orienting, and navigating present new challenges to the interface designer. In response to these changes, a body of research is developing which is beginning to address some of these interaction issues, e.g., see Crow and Pizer [1986], Chen et al. [1988], Ware [1990], Zhai and Milgram [1993], Jacob et al. [1994], Hinckley et al. [1994a], and Zhai [1995].

One of the key challenges in 3D interface design is to effectively reveal spatial relationships among objects within a 3D space, particularly in the depth dimension, so that the user can perceive, locate, and manipulate such objects with respect to each other effortlessly. This article addresses one particular 3D mechanism, the partial-occlusion effect, which can be introduced by the use of semitransparent surfaces as a means of improving 3D interaction performance. After a brief review of various 3D display techniques and the use of semitransparency in human-computer interaction (HCI), the article presents a formal experimental study of the partial-occlusion effect in a 3D interaction task. The experimental results are discussed with particular emphasis on the semitransparency characteristics and the modeling of multiple depth cues. Finally, some future research issues and potential applications of the interactive semitransparency effect are proposed.

2. BACKGROUND

2.1 Depth Cues and 3D Display Techniques

A variety of techniques are commonly used in computer interfaces for presenting 3D information. Almost all of these techniques can be linked to the depth cues identified in psychological research on human perception in the natural environment (see Haber and Hershenson [1973], Kaufman [1974], Wickens et al. [1989], and McAllister [1993] for reviews of depth cue theory). The most commonly exploited depth cues include occlusion, perspective, shadows, texture, binocular disparity, motion parallax, and active movement.

Occlusion (also called interposition) is one of the most dominant cues in depth perception. Its importance has long been recognized in 3D computer graphics, commonly through the use of hidden-line/surface removal techniques. *Stereopsis*, produced from binocular disparity, is also a strong depth cue, particularly when the perceived objects are relatively close to the viewer [Yeh 1993]. Various techniques have been devised to create a stereoscopic display on a flat computer screen [Arditi 1986; McAllister 1993]. *Perspective and relative size* are also very effective cues in 3D graphics, particularly when the displayed scene has parallel lines, as noted in Brooks [1988]. *Shadows* have also been exploited for 3D interaction [Herndon et al. 1992]. Sollenberger and Milgram [1993] showed the usefulness of the effect of *motion parallax* in graphically visualizing the connectivity of complex structures such as blood vessels in the brain. Smets [1992] and Overbeeke and Stratmann [1988] demonstrated the advantages of a *movement cue* introduced by an observer's own head movement. Similarly, Arthur et al. [1993] and Ware and Arthur [1993] found that while subjects' task completion times with a head-tracking display (to introduce movement cue) and a stereoscopic setup were similar, their error rates were significantly lower with the head-tracking condition.

The relative strengths of various depth cues have also been studied. In one early cue conflict study, Schriever [1925] compared the relative influences of binocular disparity, perspective, shading, and occlusion and showed, among other things, the dominance of occlusion over disparity information. Edge-occlusion domination was also reported in Braunstein et al. [1986]. More recently, Wickens et al. [1989], in a review of the depth combination literature, concluded that motion, disparity, and occlusion are the most powerful depth cues for computer displays.

As we can see, many of the depth cues have been carefully investigated and consciously applied to graphical displays. We noticed that yet another phenomenon—partial occlusion—produced by *semitransparent*[1] surfaces can be also a strong depth cue. Whenever a semitransparent surface overlaps another object, the viewer will see the overlapped object in lower contrast (*partially occluded*) (Figure 1). A typical example of this phenomenon in everyday life is the silk stocking; hence we also refer to the partial-occlusion phenomenon as the "silk" effect.

The partial-occlusion effect due to semitransparency is closely related to the total-occlusion cue. Although occlusion is a dominant cue in depth perception, it is often difficult to use in 3D interaction tasks, because distal objects are completely obscured by the proximal, opaque surface, leaving the user uncertain about what objects are in the background. A semitransparent surface, on the other hand, allows the user to see objects both in front and behind it. The research question here is whether partial occlusion is still a depth cue that can be readily perceived by the viewer. In other words, can viewers easily comprehend the depth relation between a semitransparent surface and other objects that are in front of or behind it? Answers to such questions are not readily available in the literature,

---

[1] A related term, "translucent," could conceivably have been used here. However, we avoid that term in this article due to our concern that it could also be construed as meaning "transmitting and *diffusing* light so that objects beyond cannot be seen clearly" (Merriam Webster's Collegiate Dictionary, 10th Edition, emphasis added).

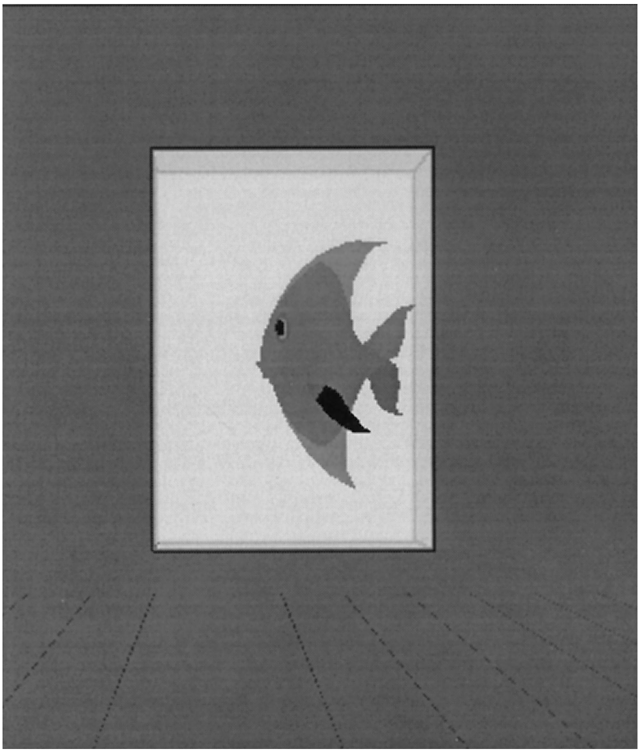

Fig. 1. Portions of an object appearing in front of (the protruding fin) or behind the semitransparent "silk" surface are perceived as such according to different levels of contrast.

possibly because semitransparency is not experienced very commonly in the natural environment. In practice, visual artists have long made use of semitransparency to enhance 3D effects in their graphical designs. With a little attention, we can see examples of semitransparency applications in television broadcasting almost every day, particularly in leading graphics of a program. The following subsection reviews a few examples in HCI designs that take advantage of semitransparency.

2.2 Examples of Semitransparent Interfaces

2.2.1 *Information Visualization.* A number of interactive systems for information visualization make use of semitransparency in rendering 3D objects. Two examples of this are the "cone tree" (Figure 2) [Robertson et al. 1991] and the "spiral calendar" (Figure 3) [Card et al. 1994]. In the case of the cone tree, as the authors state, "The body of each cone is shaded transparently, so that the cone is easily perceived yet does not block the view of cones behind." In addition, the different contrast ratio of the semitransparent cones also provides cues regarding the interrelationships of the cones in the depth dimension. In the case of the "spiral calendar" the use of semitransparent surfaces helps the user to perceive the spatial relationships among the different calendar cards. For example, the card for

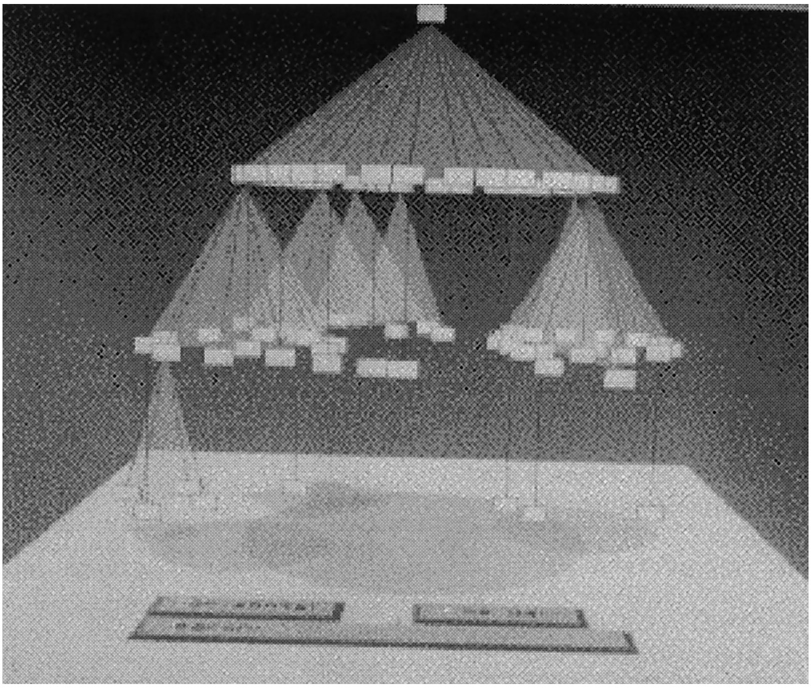

Fig. 2. The Cone Tree: The semitransparent cone bodies reveal spatial interrelationships in the depth dimension. Reprinted from Robertson et al. [1991].

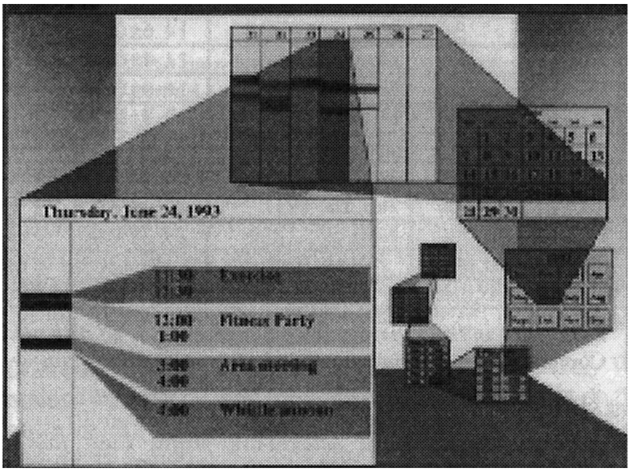

Fig. 3. The Spiral Calendar: The semitransparent surface improves the spatial structure of the interface. Reprinted from Card et al. [1994].

"June" appears to be in front of the card for "1993," due to the use of size, perspective, and partial occlusion cues.

2.2.2 *Surgical Visualization.* In a system developed for neurosurgical visualization, Hinckley et al. [1994b] used a semitransparent graphical

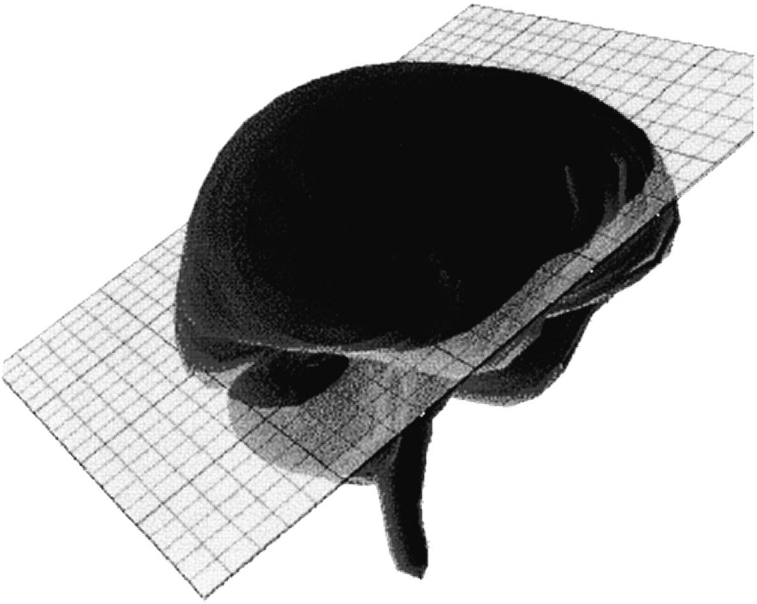

Fig. 4. A semitransparent cutting plane for surgical planning, enabling the user to see parts of the organ in front of and behind the cutting plane. Reprinted from Hinckley et al. [1994b].

representation of a cutting plane to enable surgeons to clearly visualize the spatial relationships among portions of the planned surgical procedure (Figure 4). Note that the surgeon can see parts of the organ both in front of and behind the cutting plane.

2.2.3 *Six-DOF Pursuit Tracking.* Pursuit tracking is an effective experimental research paradigm for the study of human motor skills [Poulton 1974]. It also encompasses important elements of many practical tasks, such as those found in space teleoperation. Tracking a target in six-degrees-of-freedom (6-DOF) motion requires a well-designed display to reveal mismatches (tracking errors) between the pursued target and the tracking cursor. Zhai and Milgram [1994] designed a 3D cursor with semitransparent surfaces for a 6-DOF tracking experiment that required subjects to follow and capture a target moving randomly in 6 DOF. As shown in Figure 5 the partial-occlusion cue helps the user to detect both translational (Figure 5(a)) and orientational (Figure 5(b)) errors between the target and the cursor.

2.2.4 *Two-and-a-Half-Dimensional Interfaces.* Semitransparency has also been used in 2D, multilayered interfaces (2.5D interfaces). One example of this is the “tool glass” by Bier et al. [1993] and Kabbash et al. [1994]. Figure 6 shows one example of a tool glass, in which a user can move the color palette over an object and simultaneously view both the object to be colored and the colors available for selection.

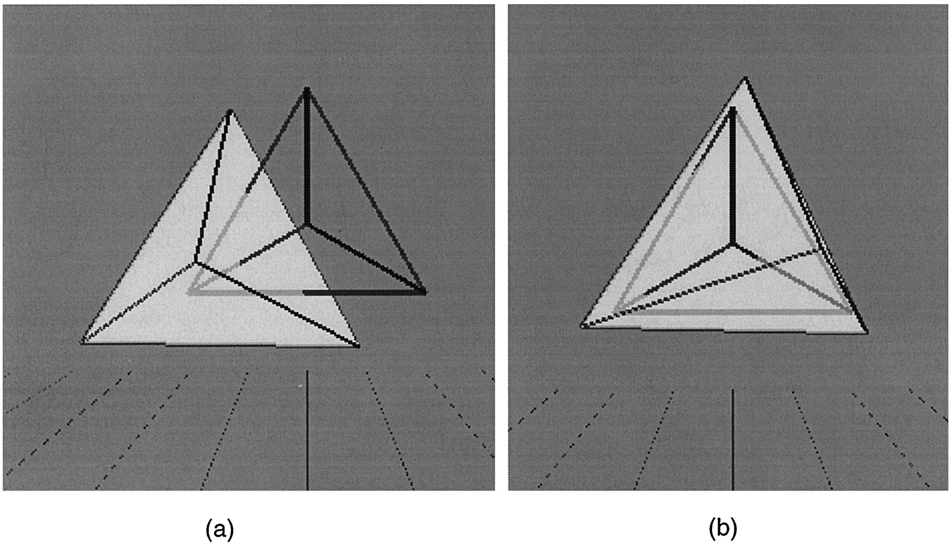

Fig. 5. Tracking a 3D target with a "silk cursor": Translational (a) and rotational (b) differences between the cursor and the target are effectively revealed with the silk surfaces.

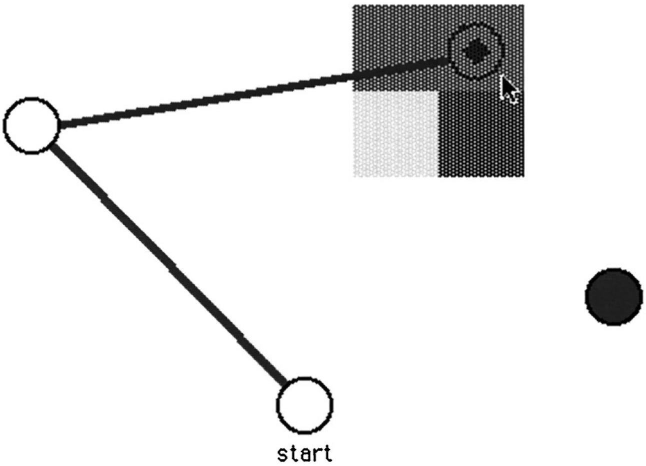

Fig. 6. Color selection "tool glass": The user can superimpose the *semitransparent* color plate on a target object and click *through* the color selected for drawing. Courtesy of Paul Kabash.

Another potential application of semitransparency is with user interface (UI) widgets such as pull-down menus and dialogue boxes. Conventional widgets often obscure the very objects on which the user wishes to focus attention. One way to solve this problem could be to use a semitransparent background when constructing the UI widget so that the user can manipulate the widget while still seeing the objects underneath. Figure 7 shows a snapshot of *SilkWidgets*, a sketching program developed by the authors at Alias|Wavfront to test the concept of semitransparent widgets. One issue in applying semitransparent widgets is the possible interference between

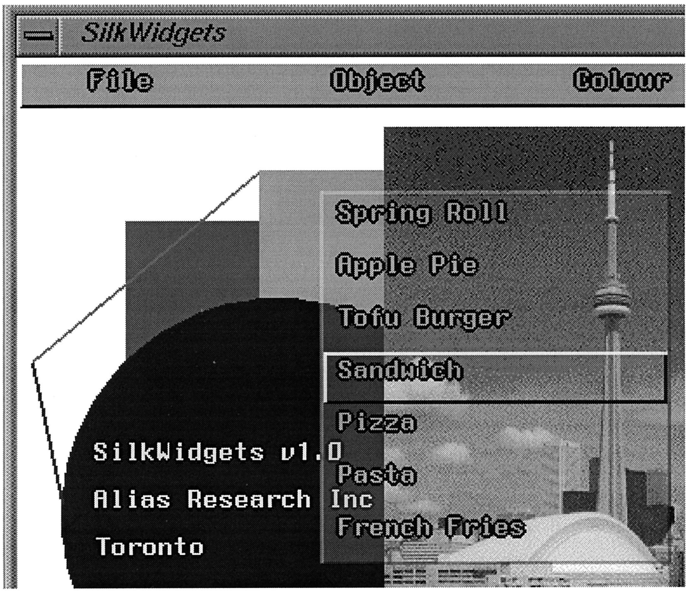

Fig. 7. Semitransparent popup menu in SilkWidgets. Copyright Alias Research, Inc.

information contained in the widgets and the objects underneath. This has been addressed in Harrison et al. [1994; 1995].

2.2.5 *Using Transparency as State Display.* Semitransparency has also been used in ways other than introducing depth cues per se to display "state" in 3D interaction. Venolia [1993] used semitransparency as interactive feedback for the state of 3D objects. Whenever an object was "touched," i.e., a small cursor moved into the object, the object changed from opaque to semitransparent. The user could see not only the touched object, but also the cursor inside the object.

Many more examples of semitransparency in human-computer interfaces can be found. Despite the numerous applications, however, the effectiveness of using semitransparency in human-computer interaction has seldom been studied formally. We hypothesize that viewers can perceive the position of a semitransparent surface *in relation to* other objects in 3D environments. Furthermore, we believe that the partial-occlusion cue through semitransparency is particularly useful in 3D *interaction*, because as users gradually move a semitransparent surface *through* an object, they can perceive the immediate and continuous change in the object's appearance. This suggests a potentially powerful mechanism for users to locate objects in 3D interaction tasks. We carried out a quantitative experimental study, in order to determine the power, the characteristics, and the limitations and constraints of semitransparency in 3D interaction relative to other commonly used 3D techniques such as stereoscopic presentation.

It is also important to note that semitransparency is relatively easy to implement with today's computer systems, which further increases the

Fig. 8. The experimental setup.

justification for its careful study and wider application in computer interfaces.

3. METHOD

3.1 Experimental Task

In each trial of the experiment, a graphically rendered angel fish moved randomly in x, y, and z dimensions within a 3D virtual environment (Figure 8). Subjects were asked to manipulate a 3D cursor (Figure 9), with or without the silk surface, to envelop the fish and to "grasp" it when the fish was perceived to be completely inside the cursor. Subjects wore a special glove (Figures 8 and 10) as the input device (Section 3.1.3). Grasping was done simply by closing the hand naturally. If the fish was entirely inside the cursor volume the trial was successful, and the fish thus stayed "caught." The time score of each trial was displayed to the subjects, along with a short beep. If the fish was not *completely* inside the cursor when grasped, the fish disappeared. In this case (considered a "miss") a long beep was sounded, and error magnitudes in each of the x, y, and z dimensions were displayed, along with the message "Missed." Each new trial was activated when the subjects pressed the spacebar on the keyboard. Subjects were instructed to complete each trial as quickly as possible and catch as many fish as possible.

Although presented as a game (which was greatly enjoyed by the subjects), the virtual "fishing" task is essentially a 3D dynamic target acquisition task, comprising both perception and manipulation in 3-space. Note that in this study target acquisition was taken as an *experimental scenario* to test user performance in perceiving and positioning objects in 3D, which are often fundamental elements in many of the 3D interaction tasks, such

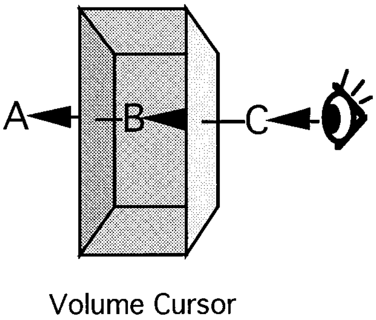

Fig. 9. Use of a "silk" covering over a rectangular volume cursor in order to obtain occlusion-based depth cues. An object at point A is seen through two layers of "silk" and is thus perceived to be *behind* the cursor. An object at point B is seen through one layer and thus is perceived as being *inside* the cursor's volume. An object at point C is not occluded by the silk at all and so is seen to be *in front of* the volume cursor.

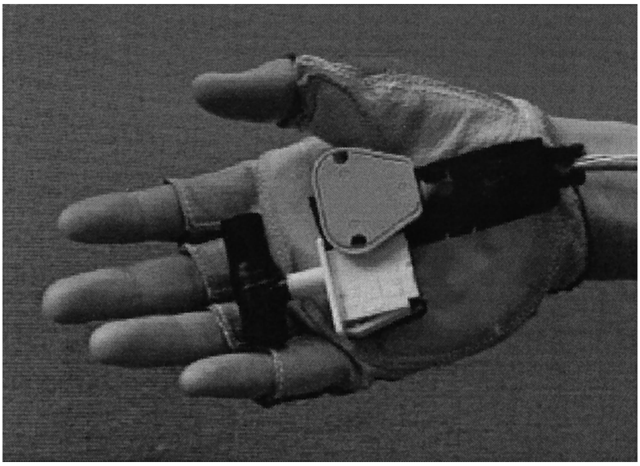

Fig. 10. The input glove.

as acquiring objects, moving, dragging, rotating, stretching or sculpting them, or navigating along a desired trajectory. The silk cursor is not necessarily a practical 3D *target selection technique* in the narrow sense. Designating or selecting 3D targets as a practical task does not necessarily require much depth information, and there are simple and effective techniques for that purpose. For instance, the subject can easily "shoot" at a graphical fish with a line of ray trace or a virtual "spotlight" as described in Liang and Green [1994].

3.1.1 *Experimental Platform.* The experiment was conducted using MITS (Manipulation In Three Space), a desktop stereoscopic virtual environment developed by the authors. The experiment described here was carried out on an SGI IRIS Crimson/VGX graphics workstation. The MITS system automatically records a broad range of information during the experiment and manages the timing and execution of the experiment, including presentation of instructions to subjects so that experiments can be run with minimal interference or bias from the experimenter.

The origin of the {x, y, z} coordinates of the MITS virtual environment was located at the center of the computer screen surface, with the positive x-axis pointing to the right; the y-axis pointing upward, and the z-axis pointing toward the viewer. All objects were drawn using perspective projection and were modeled in units of centimeters, where 1 cm in the virtual fish tank corresponded to 1 cm in the real world for any line segment appearing within the same plane as the surface of the screen. The graphics update rate was controlled at 15Hz in this experiment.

3.1.2 *The Targets and Their Motion.* Each of the targets ("angel fish") used in this experiment had a flat body, except for two fins and two eyes protruding from the body (Figure 1). The angle between any fin and the body was 30 degrees. The size of the fish varied from trial to trial. The x (from lips to tail), y (vertical), and z (from left fin tip to right fin tip) dimensions of the largest ("adult") fish were 10 cm, 15 cm, and 1.3 cm, respectively. The smallest ("baby") fish was 30% the size of the largest adult fish.

The fish movements were driven by independent forcing functions in the x, y, and z dimensions. Such inputs, based on suitable combinations of sinusoidal functions, generate smooth and subjectively unpredictable motion and are employed quite frequently in manual-tracking research [Poulton 1974]. In this experiment, the particular forcing functions applied to the fish motions were

$$x(t) = \sum_{i=0}^{5} Ap^{-i} \sin(2\pi f_0 p^i t + \phi_x(i))$$

$$y(t) = \sum_{i=0}^{5} Ap^{-i} \sin(2\pi f_0 p^i t + \phi_y(i))$$

$$z(t) = -7.8 + \sum_{i=0}^{5} Ap^{-i} \sin(2\pi f_0 p^i t + \phi_z(i))$$

where t was the time from the beginning of each test (see Section 3.3 on experimental design and procedure for the definition of a test), $A = 4.55$ cm, $p = 2$, and $f_o = 0.02$Hz. The phase terms, $\phi_x(i)$, $\phi_y(i)$, and $\phi_z(i)$ ($i = 0, 1, \ldots 5$), were pseudorandom numbers, ranging uniformly between 0 and 2π. This design resulted in fish motions which were sufficiently unpredictable to the subjects and different from trial to trial, but repeatable for each test and between experimental conditions.

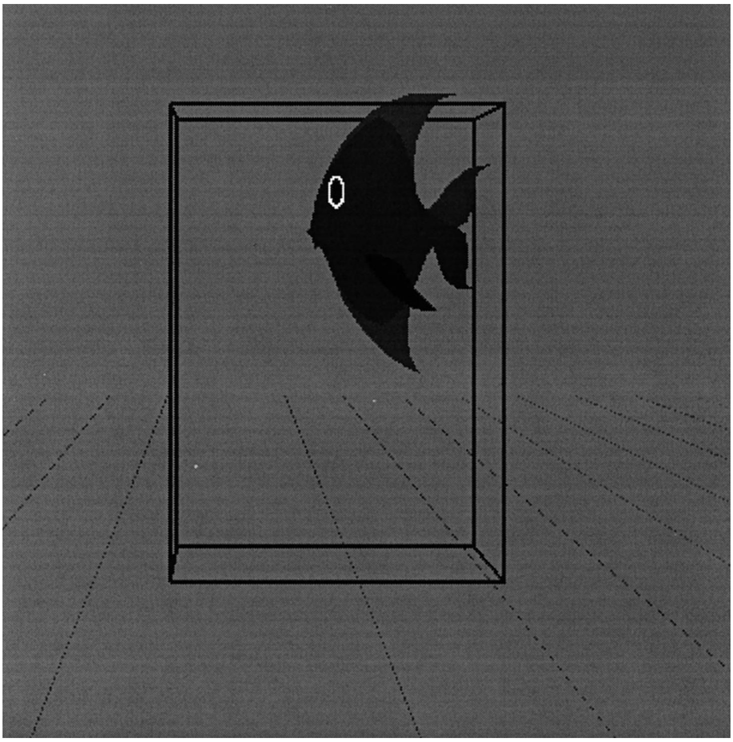

Fig. 11. A fish and the wireframe cursor.

3.1.3 *The Cursor and the Input.* The cursor used to capture the fish was a rectangular box of size 11.3 cm, 16.3 cm, and 2.6 cm in x, y, and z dimensions, respectively (Figure 9). Two versions of the cursor were used in the experiment. One was a *wireframe cursor* that had no surfaces (Figure 11), and the other was a *silk cursor* with semitransparent surfaces (Figure 1, Figures 12–14). The intensity, I, of the semitransparent surface was rendered by blending the cursor color (source) intensity, I_S, with the destination color intensity, I_d, [Foley et al. 1990], according to

$$I = \alpha I_s + (1 - \alpha) I_d.$$

Although I_S was chosen to be white in this experiment, different color compositions may be more suitable for other particular applications.

If $\alpha = 1$, the cursor is totally opaque and therefore completely occludes objects behind it. If $\alpha = 0$ the cursor is totally transparent, and no partial occlusion cues are available. On the basis of pilot experiments, we determined a suitable coefficient of $\alpha = 0.38$ for all surfaces of the cursor, except for the back surface, which was set at $\alpha = 0.6$. These values resulted in partial-occlusion states (i.e., in front, between two layers, and behind two layers of silk surface) that were judged to be satisfactorily distinguishable.

The wireframe cursor as used in the experiment (Figure 11) can obviously be improved by drawing line segments or cross hairs on the cursor

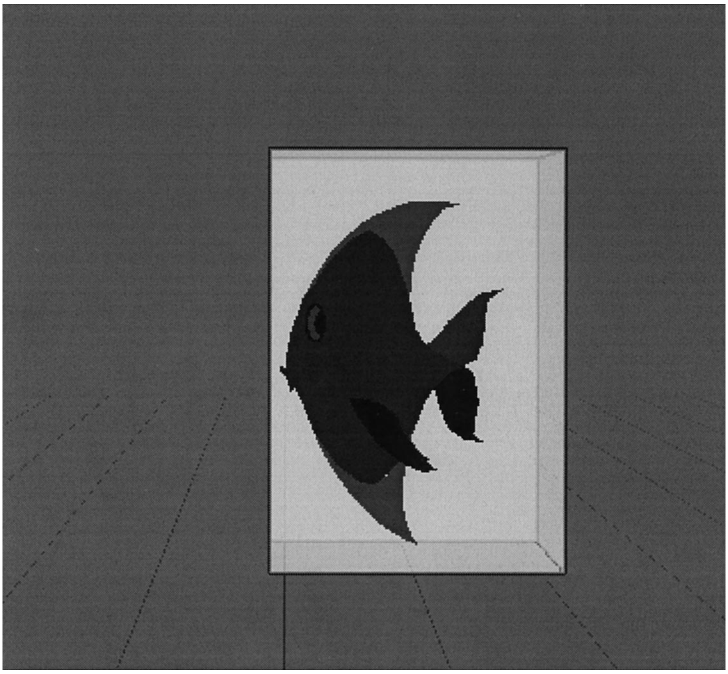

Fig. 12. A fish *in front of* the silk cursor.

surface so that the cursor appears like a fishing net. The resulting effect is also a form of "partial occlusion." When the fishing net is dense enough, it will appear semitransparent. In fact, this is one of the approaches, often called the "screen door" approach, for implementing semitransparency in computer graphics [Foley et al. 1990, p. 755]. In order to investigate the effect of partial occlusion, we choose two special levels of occlusion: the wireframe cursor represents the extreme case of no occlusion (except the cursor edges), while the silk cursor exhibits an optimized degree of partial occlusion.

In the experiment, the cursor was driven by a custom-designed glove based on an Ascension Technology Bird™ equipped with a clutch, as shown in Figures 8 and 10. The glove operated in position control mode, with a Control/Display ratio of 1:1, as determined in previous research [Zhai 1995]. The "home" positions of the glove corresponded to a cursor location of (0, 0, 0) and were calibrated to make the subject most comfortable when using the glove. The Bird receiver and the clutch were at the center of the user's hand to best allow the user to "grasp" a fish by means of finger/hand abduction. The Bird has six degrees of freedom (x, y, z, roll, pitch, yaw). Since only translations were needed in the task, rotational signals were disabled for this experiment.

3.1.4 *The Display.* The fishing task was displayed on an SGI monitor with a resolution of 1280 by 1024 pixels (model no. HL7965KW-SG). Monitor brightness and contrast were adjusted so as to minimize ghosting

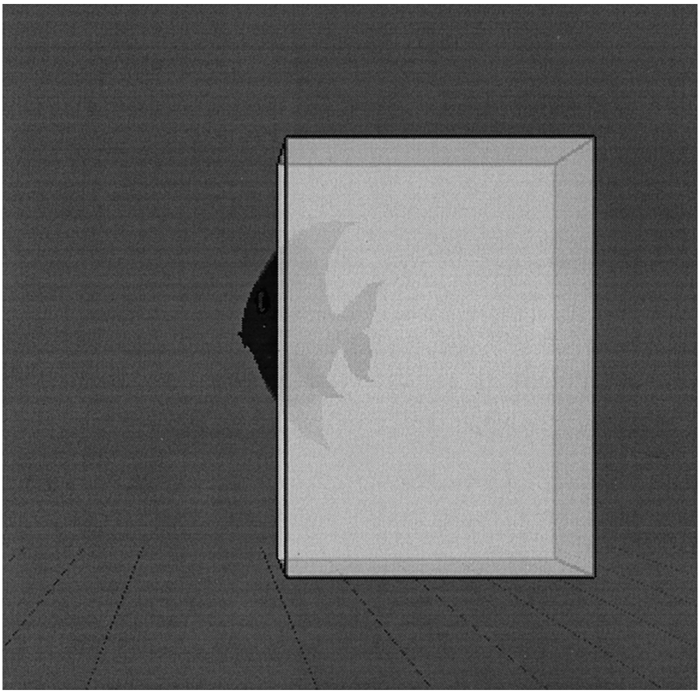

Fig. 13. A fish *behind* the silk cursor.

images for the stereoscopic displays and thereby optimize the stereoscopic effect. The experimental room was darkened throughout the experiment. The gamma correction value was set at 1.70.

Two modes of display were used in the experiment: stereoscopic and monoscopic projection. In the stereoscopic case, subjects wore 120Hz flicker-free stereoscopic CrystalEyes™ glasses (model no. CE-1), manufactured by StereoGraphics.

3.2 Experimental Conditions and Hypotheses

The primary goal of our experiment was to evaluate the effectiveness of the partial occlusion cue in 3D interaction. Stereoscopic projection has been found to be a powerful technique in displaying depth information [McAllister 1993; Wickens et al. 1989; Yeh and Silverstein 1992] and has been used often as the control condition in studying other types of 3D display techniques; e.g., Sollenberger and Milgram [1993] and Arthur et al. [1993]. This experiment was designed to allow comparison of the relative performance of stereopsis versus partial occlusion for the interaction task. We were also interested in learning how the two sources of cues interact. Thus, two display modes (monoscopic versus stereoscopic) and two types of cursor (silk cursor versus wireframe cursor) were included in the experiment, resulting in four conditions: silk cursor with stereo display (SilkStereo), wire frame cursor with stereo display (WireframeStereo), silk cursor with

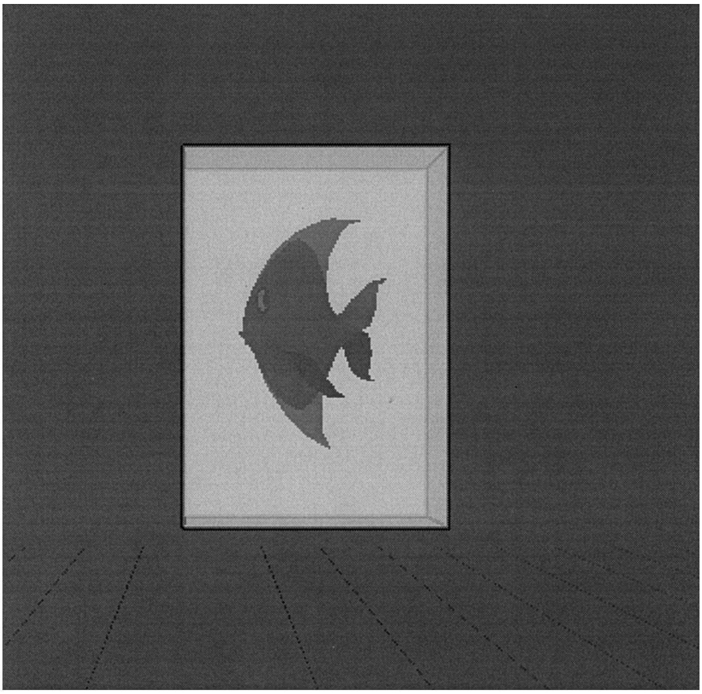

Fig. 14. A fish completely *inside of* the cursor.

mono display (SilkMono), and wire frame cursor with mono display (WireframeMono).

Apparently, the WireframeMono case, the baseline condition, is the most difficult one since neither partial occlusion nor stereopsis was present for judging depth relation. The subjects had to rely on occlusions between the edge of the cursor and the fish. They tended to move the cursor so that the fish first was apparently located between the edges of the cursor in the z dimension (Figure 11) and then slightly adjust the cursor in the x and y dimensions to bring the fish into the center of the cursor before grasping.

In the WireframeStereo case, subjects no longer had to depend on edge occlusion. Because the stereoscopic cue gave them a strong 3D sensation, they could judge the depth dimension directly and simultaneously with their judgment along the and x and y dimensions.

In the SilkMono case, portions of the target appeared with different contrast ratios when they were located in front of (Figure 12), behind (Figure 13), or inside the cursor (Figure 14). The subjects tended to use the semitransparency cue *interactively*, by moving the silk cursor first *through* the target to observe the continuous change of target appearance (see Figure 1, the portion of fin in high contrast will change as the cursor moves) and then grasping *immediately* after the front surface of the silk cursor moved in front of the fish fin. The ability to judge where the semitransparent surface is relative to the target through interactive movement is critical to the power of the partial-occlusion cue. Without this interactive effect, subjects would not be able to tell when the back fin of the fish was inside the cursor (see Figure 14).

In the SilkStereo case, subjects had the advantage of both the stereo cue and the partial-occlusion cue. We expected SilkStereo to be the most efficient case and WireframeMono to be the least efficient. Whether SilkStereo would be significantly superior to WireframeStereo would reveal whether the partial-occlusion cue provides depth information in addition to stereo cue. What was also of particular interest to us was whether the SilkMono case (partial occlusion cue alone) would generate superior, or in any case comparable, performance scores relative to the case of WireframeStereo (stereo cue alone), which would confirm to us the potentially powerful advantages of the semitransparency on its own.

Stated formally, our hypotheses for this particular class of localization tasks were

(1) partial occlusion improves performance;
(2) stereoscopic display also improves performance;
(3) the strength of the partial-occlusion cue is no less powerful than the stereo cue; and
(4) the two cues enhance each other, and performance is best when both cues are present.

3.3 Experimental Design and Procedure

Twelve paid subjects were recruited through advertising on campus. The subjects were screened using the Bausch and Lomb Orthorator visual acuity and stereopsis tests. Subjects' ages ranged from 18 to 36, with the majority in their early and mid-20's. One of the 12 subjects was left-handed and the rest were right-handed, as determined by the Edinburgh inventory [Oldfield 1971]. Subjects were asked to wear an input glove on their dominant hand.

A balanced within-subjects design was used. The 12 subjects were randomly assigned to an order of the four conditions (SilkStereo, WireframeStereo, SilkMono, WireframeMono) using a hyper-Graeco-Latin square pattern, which resulted in every condition being presented an equal number of times as a first, second, third, and final condition.

Following a two-minute demonstration of all four experimental conditions, the experiments with each subject were divided into four *sessions*, with one experimental condition in each session. There was a one-minute rest period between every two sessions. Each session comprised five *tests*. Each test consisted of 15 *trials* of fish catching. Test 1 started when the subject had no experience with the particular experimental condition. Tests 2, 3, 4, and 5 started after the subjects had 3, 6, 9, and 12 minutes worth of experience respectively. Practice trials filled the gap following a test and before the next test began, so that each test (e.g., Test 3) always started when the subject had a fixed amount of practice with the particular experimental condition (e.g., nine minutes for Test 3). At the end of each test, the number of fish caught and missed (as both an absolute number and a relative percentage) and mean trial completion time were displayed to the subject.

At the end of the experiment, a short questionnaire was administered to assess users' subjective preferences for all experimental conditions.

3.4 Performance Measures

Task performance was measured by trial completion time, error rate, and error magnitude. Trial completion time was defined as the time duration from the beginning of the trial to the moment when the subject grasped. Error rate was defined as the percentage of fish missed in a *test* (15 trials). Whenever a fish was missed, the error magnitude was defined as the Euclidean summation of errors (portions of the body *outside* of the cursor) in the x, y, z dimensions, respectively:

$$\text{Error Magnitude} = \sqrt{e_{\mathrm{x}}^2 + e_{\mathrm{y}}^2 + e_z^2}$$

Note that the error magnitude is not a primary measure for two reasons. First, the subjects' task was to capture the fish as quickly as possible. Error magnitude was not an explicit requirement. Second, it occurred only when the subject missed the fish. We included error magnitude to gather a complete set of performance measures.

3.5 Experimental Results

Three thousand and six hundred experimental trials (i.e., 12 (subjects) × 2 (cursor types) × 2 (display modes) × 5 (tests) × 15 (trials per test)) of data were collected during the experiment. Repeated-measure analyses of variance were conducted through the multivariate approach [Bock 1975] to test the statistical significance of the individual effects and their interactions under each of the three performance measures. All of the analyses were conducted first with the original performance score, followed by examination of the variances of the different effects and model residuals. As is common in human subject experiments, the data on trial completion times, error rates, and error magnitudes collected here were not normally distributed, but rather skewed toward lower values. In order to increase the validity of the statistical analysis [Howell 1992], logarithmic transformations were applied to the trial completion time and error magnitude data, and a square root transformation was applied to the error rate data. These transformations made the data meet the variance analysis assumptions of normality and homogeneity of variance. Statistical results (Section 3.5.1) were based on the transformed data, even though graphs are presented in the original scale for ease of comprehension. Greenhouse-Geisser and Hunyh-Feldt adjustment epsilon values were calculated to estimate the potential correlation in repeated-measure designs, but no critical differences were found between the original and the adjusted probability values.

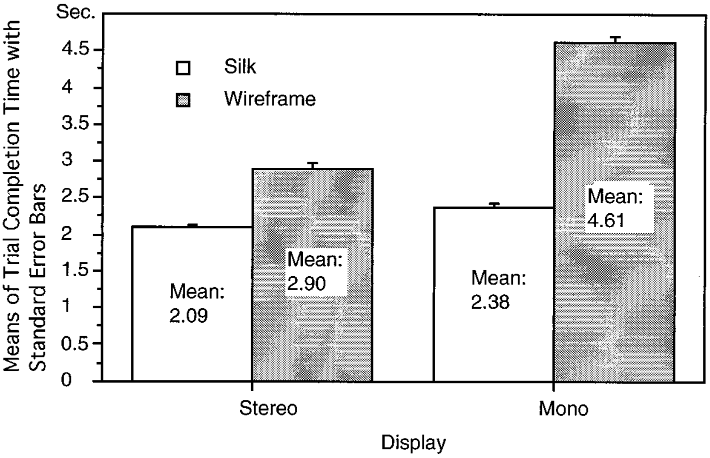

Fig. 15. Trial completion times as a function of cursor type and display mode.

The reported probabilities were therefore left unadjusted. The following are the primary results of the statistical analysis.

3.5.1 *Trial Completion Time.* Variance analysis indicated that cursor type (silk vs. wireframe cursor: $F_{1,11} = 66.47$, $p < 0.0001$), display mode (stereo vs. mono display: $F_{1,11} = 15.0$, $p < 0.005$), experimental phase ($F_{4,44} = 21.59$, $p < 0.0001$), trial number (different fish size and 3D location: $F_{14,154} = 12.55$, $p < 0.0001$), cursor × display interaction ($F_{1,11} = 6.68$, $p < 0.05$), and cursor × display × phase interaction ($F_{4,44} = 4.0$, $p < 0.01$) all significantly affected trial completion time.

Figure 15 illustrates the effect of cursor type and display mode on trial completion time. Multiple contrast tests showed that the silk cursor produced significantly shorter completion times than the wireframe cursor for *both* monoscopic and stereoscopic displays (Table I). With regards to the magnitude of the differences, the mean completion time with the silk cursor was 48.4% shorter than that with the wireframe cursor in monoscopic display and 28.1% shorter in stereoscopic display. Finally, the mean completion time for SilkMono (partial-occlusion cue alone) was 18.1% shorter than for WireframeStereo (stereo cue alone), even though this difference was not statistically significant ($p = 0.28$). These results suggest that, under the experimental conditions, (1) the use of semitransparent surfaces brought significant benefit to task performance as measured by completion time and (2) the power of partial occlusion through semitransparency was comparable, if not stronger than that of stereopsis.

3.5.2 *Error Rate.* As illustrated in Figure 16, the pattern of the error rate data as a function of cursor type and display mode is very similar to that of the trial completion time data. The statistically significant factors affecting error rate were cursor type ($F_{1,11} = 92.16$, $p < 0.0001$), display mode ($F_{1,11} = 14.48$, $p < 0.01$), and cursor type × display mode interaction

Table I. Multiple Contrast Tests of Mean Completion Times

	P-Value
SilkStereo vs. SilkMono	0.31
SilkStereo vs. WireframeStereo	0.05
SilkStereo vs. WireframeMono	0.0001
SilkMono vs. WireframeStereo	0.28
SilkMono vs. WireframeMono	0.0001
WireframeStereo vs. WireframeMono	0.0006

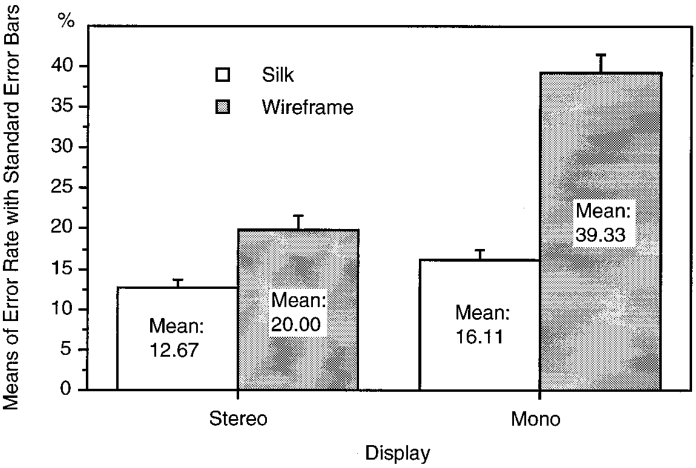

Fig. 16. Error rate as a function of cursor type and display mode.

($F_{1,11}$ = 7.47, $p < 0.05$). Neither experimental phase nor any interactions between experimental phase and other factors were significant.

Multiple contrast tests showed that the silk cursor produced significantly fewer errors than the wireframe cursor, both for monoscopic displays and for stereoscopic displays (Table II). Regarding the actual differences in magnitude, for monoscopic displays the mean error rate of the silk cursor was 59% less than that of the wireframe cursor. For stereoscopic displays the mean error rate with the silk cursor condition was 36.7% less than for the wireframe cursor. For the case of partial-occlusion cue alone (SilkMono), the mean error rate was 19.5% lower than for the stereo cue alone (WireframeStereo), but this difference was not statistically significant (p = 0.21). Similar to the trial completion time data, the error rate data suggest that the partial-occlusion cue indeed brought performance improvement relative to the control condition, and it was no less powerful than the stereopsis cue.

3.5.3 *Error Magnitude.* The effects of cursor type and display mode on error magnitude are shown in Figure 17. When examining the error magnitude data, please bear in mind that error magnitude was defined only when an error was made (i.e., a target was missed) and that fewer errors occurred in some conditions than for others, as indicated. The

Table II. Multiple Comparison Tests of Mean Error Rate

	P-Value
SilkStereo vs. SilkMono	0.21
SilkStereo vs. WireframeStereo	0.02
SilkStereo vs. WireframeMono	0.0001
SilkMono vs. WireframeStereo	0.21
SilkMono vs. WireframeMono	0.0001
WireframeStereo vs. WireframeMono	0.0003

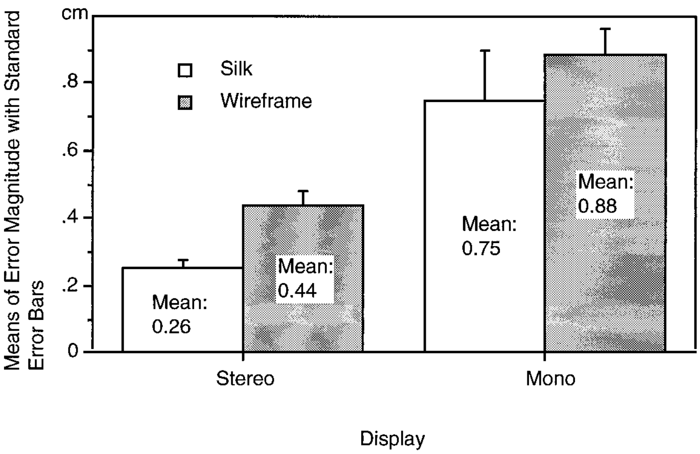

Fig. 17. Error magnitude as a function of cursor type and display mode.

variance analysis concluded that error magnitude was significantly affected by cursor type ($F_{1,11} = 11.37$, $p < 0.01$), display mode ($F_{1,11} = 18.19$, $p < 0.001$), and experimental phase ($F_{4,44} = 3.97$, $p < 0.01$). No significant between-factors interactions of any order were found.

Multiple contrast tests (Table III) showed that the silk cursor produced significantly lower error magnitudes than the wireframe cursor, both for monoscopic displays and for stereoscopic displays. For monoscopic displays the mean error magnitude of the silk cursor was 15.1% smaller than that of the wireframe cursor. For stereoscopic displays the mean error magnitude of the silk cursor condition was 41.5% smaller than that of the wireframe cursor.

In contrast to the trial completion time and error rate data, it appears that when an error did occur, the stereo cue was more effective than the partial-occlusion cue in reducing the error magnitude. The SilkMono mode (partial-occlusion cue alone) produced a larger mean error magnitude average (as well as larger deviation, Figure 17) than the WireframeStereo mode (stereo cue alone). However this difference was not statistically significant ($p = 0.97$).

3.5.4 *Learning Effects and Final-Phase Results.* As indicated in the variance analyses above, the learning phase was a significant factor for

Table III. Multiple Comparison Tests of Mean Error Magnitude

	P-Value
SilkStereo vs. SilkMono	0.03
SilkStereo vs. WireframeStereo	0.03
SilkStereo vs. WireframeMono	0.0001
SilkMono vs. WireframeStereo	0.93
SilkMono vs. WireframeMono	0.005
WireframeStereo vs. WireframeMono	0.007

trial completion time and error magnitude, but not for error rate. It also interacted significantly with cursor display combinations, as measured by trial completion time. This subsection describes the performance changes, as learning progressed, and the results in the final phase of the experiment.

Figure 18 shows trial completion time data for each technique as a function of the experimental phase. It shows clearly that the relative scores between the different conditions were ordinally consistent over all experimental phases. Subjects improved their time scores for the SilkStereo, SilkMono, and WireframeStereo modes as they gained more experience and presumably more confidence. Little improvement in completion time was evident with the WireframeMono condition, however.

Variance analysis was conducted on trial completion time data in the final experimental phase (Test 5 in Figure 18). The statistical conclusions were the same as those drawn from the overall data above: cursor type ($F_{1,11} = 90.8$, $p < 0.0001$), display mode ($F_{1,11} = 21.5$, $p < 0.001$), cursor type and display mode interaction ($F_{1,11} = 17.3$, $p < 0.005$), and trial number ($F_{14,154} = 6.4$, $p < 0.0001$) all significantly affected trial completion time. Results of the multiple contrast comparisons for the final-phase completion time data also agreed with the results from the overall data (Table I): SilkStereo vs. SilkMono ($p = 0.27$) and SilkMono vs. WireframeStereo ($p = 0.32$) were not significantly different. All other pair comparisons were significant ($p < 0.05$). Mean trial completion time reductions due to the partial-occlusion effect in the final phase are as follows. For mono displays, SilkMono (mean 2.064 seconds) was 52.8% less than WireframeMono (mean 4.376 seconds). For stereo display, SilkStereo (mean 1.850 seconds) was 20.6% less than WireframeStereo (mean 2.329 seconds).

Figure 19 presents the error rate data as a function of the experimental phase. Again the relative rank of each mode was consistent across all five phases of the experiment. Interestingly however, in contrast to the completion time data (Figure 18), the error rate for the WireframeMono condition showed the most obvious improvement over the experiment. A small amount of improvement was also found in the SilkMono condition, while essentially none was found in the SilkStereo and WireframeStereo modes. Variance analysis for the final-phase (Test 5) error rate data showed that cursor type ($F_{1,11} = 26.6$, $p < 0.0005$) and display mode ($F_{1,11} = 6.05$, $p < 0.05$) were both significant factors, but the cursor type and display mode interaction ($F_{1,11} = 1.53$, $p = 0.24$) was not significant. Multiple contrast

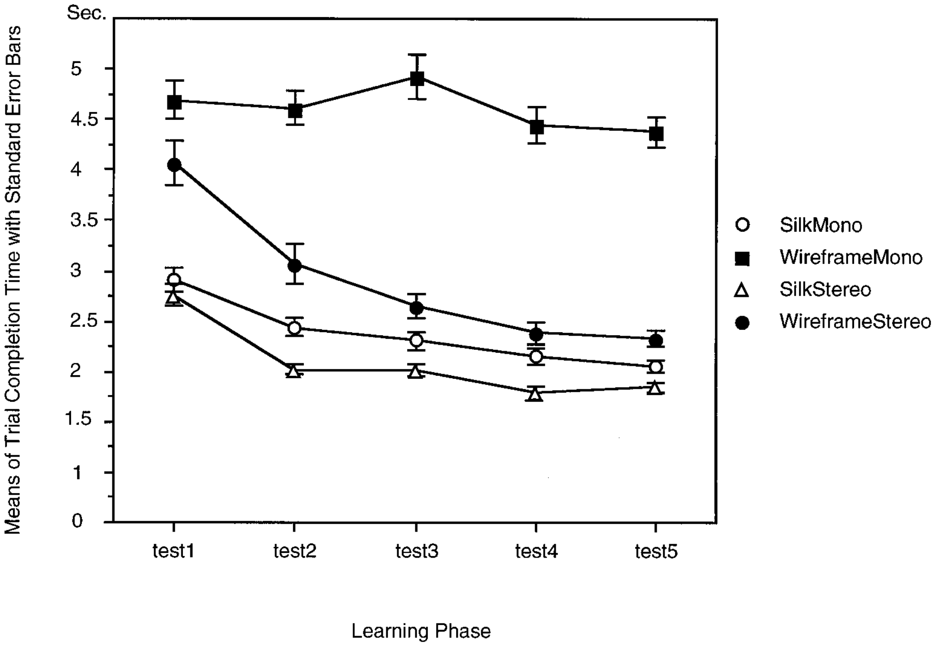

Fig. 18. Time performance for each of four conditions at each learning phase.

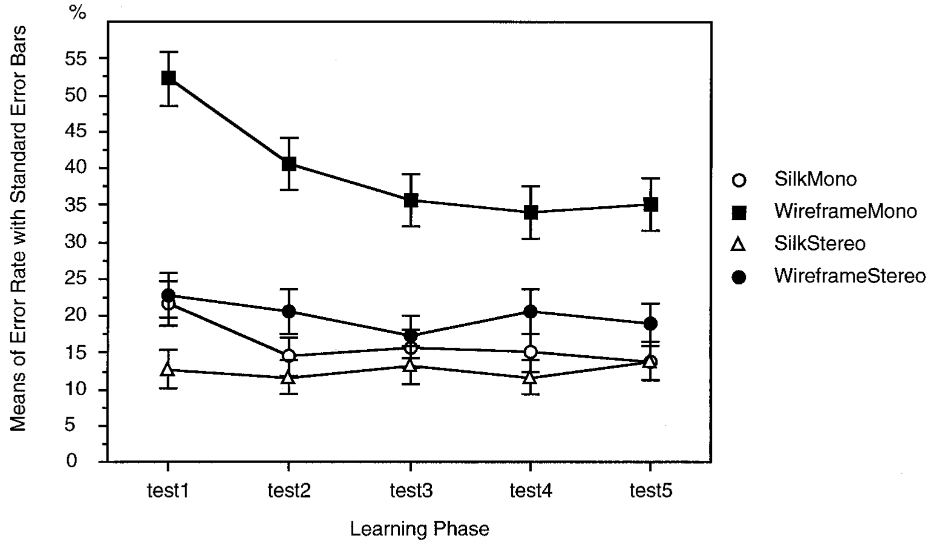

Fig. 19. Error rate for each of four conditions at each learning phase.

comparisons showed that final-phase error rate with WireframeMono was significantly higher than the other three cases ($p < 0.05$). Other contrasts were not significant ($p > 0.05$). Mean error rate reductions resulting from the partial-occlusion effect in the final phase are as follows. For the mono display, error rate with SilkMono (mean 13.9%) was 60.8% lower than

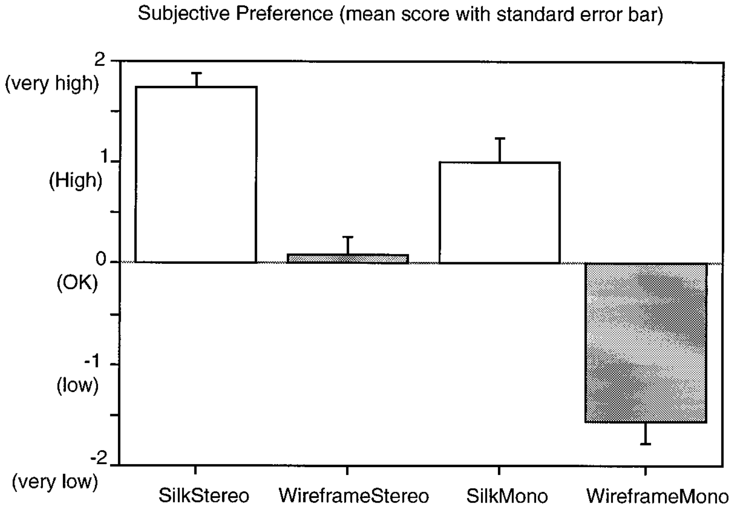

Fig. 20. Mean scores for subjective evaluation.

WireframeMono (mean 35.0%). For the stereo display, SilkStereo (mean 13.9%) was 26.5% lower than WireframeStereo (mean 18.9%). Note that the lowest average error rate (13.9%) was still greater than the error rates found in typical 2D target acquisition studies. This is probably due to two reasons. One is that the task was more difficult than usual, not only because it was performed in 3D but also because the target (fish) was always moving. The second reason is related to the instructions given to the subjects who were told to "catch as many fish as possible and complete each trial as quickly as possible." No emphasis was given to ensure that no fish was missed.

Comparing Figures 18 and 19 reveals important information about speed accuracy tradeoff patterns with respect to learning. For the WireframeMono mode, subjects had more than a 35% error rate, which apparently caused them to focus on improving the accuracy aspect of the task at the expense of time performance. In the other three cases (SilkStereo, SilkMono, and WireframeStereo), subjects already had less than a 25% error rate, and it appeared that they were more satisfied with this level of accuracy, and thus were devoting more effort to reducing their trial completion times.

The error magnitude data were not suitable for statistical analysis as a function of each experimental phase, since very few errors occurred for some of the phase and technique combinations.

3.5.5 *Subjective Preferences.* Figure 20 shows the mean scores for the subjective evaluation data collected after the experiment. On the average, SilkStereo was the most preferred and WireframeMono the least preferred, with SilkMono ranked higher than WireframeStereo. Statistically, significantly different preference scores were found across conditions through repeated-measure variance analysis ($F_{3,33} = 54.36$, $p < 0.0001$). The

Table IV. Multiple Contrast Test Results of Subjective Preference

	P-Value
SilkStereo vs. SilkMono	0.01
SilkStereo vs. WireframeStereo	0.0001
SilkStereo vs. WireframeMono	0.0001
SilkMono vs. WireframeStereo	0.002
SilkMono vs. WireframeMono	0.0001
WireframeStereo vs. WireframeMono	0.0001

results of the multiple contrast tests are summarized in Table IV and show that subjects' preferences between every pair of techniques were significantly different (including WireframeStereo vs. SilkMono). Interestingly, the subjective evaluation data in this experiment were consistent with the acquired performance measures (completion time and error rate) in trend, but were more sensitive in detecting differences between conditions.

3.5.6 *Summary of Results.* The experiment largely confirmed our initial hypotheses. In terms of all three measures of performance (trial completion time, error rate, and error magnitude), both stereopsis through binocular disparity and partial occlusion through semitransparency were significantly beneficial to the manual 3D localization task. The partial-occlusion cue was effectively used by subjects in both display modes: it significantly improved users performance not only in the monoscopic display which had little depth information available, but also in the stereoscopic display which already had the powerful stereo cue. Comparing the two cues, partial occlusion was no less powerful than stereopsis for successful 3D target acquisition. Learning improved subjects' performance with each of the techniques, but the relative rank of the techniques remained unchanged throughout the experiment. Subjective evaluations supported the conclusions drawn from performance measures.

4. DISCUSSIONS

4.1 Properties of Semitransparency: Discrete, Relational Depth Cueing

Two particular properties of semitransparency are especially relevant to practical HCI applications. One of these is the fact that a semitransparent surface does not completely block the view of any object which it (partially) occludes. This eliminates one of the *disadvantages* of the powerful total-occlusion cue and permits the user to maintain awareness of the background information.

The second property relates to the fact that, similar to the total-occlusion cue, partial occlusion through semitransparency provides *relational and discrete* depth information about the position of a semitransparent surface *relative to* other objects. This is in contrast to stereoscopic displays, which provide continuous, quantitative depth information. As illustrated in Figure 1 and Figures 12 to 14, we see how the silk covering the volume cursor directly reveals *whether* an object is in front of the cursor, within it, or behind it. When an object is behind a semitransparent surface, however, the user will not able to tell by *how much* the object is separated from the surface in space. For some tasks, such as making an absolute judgment of distance, the discrete nature of the partial-occlusion cue may represent a shortcoming, whereas for others it will be a distinct advantage, since the user does not have to make a qualitative decision based on quantitative, continuous information. This was precisely the case in the experiment described here, where the objective was to manipulate the cursor so that it totally enveloped the fish being hunted. This is clearly a discrete task, as the subjects were instructed simply to capture the fish and not necessarily to center the cursor on it as accurately as possible. This contention is supported by evidence from the experiment: in Figures 15 and 16 we see that semitransparency appears to be a slightly more effective cue than binocular disparity for successful target acquisition. However, upon examining Figure 17, we note that the mean error magnitude and variance of the SilkMono case were larger than those of the WireframeStereo case. The implication of this is that although fewer errors were made under the SilkMono condition relative to the WireframeStereo condition, the magnitude of those fewer errors must have been relatively larger than in the WireframeStereo case, suggesting the distinction between discrete and continuous depth information.

We also point out that, although static semitransparent surfaces provide primarily discrete cues, continuous depth information can nevertheless be acquired when semitransparent surfaces are used as a *dynamic interactive medium*. That is, when the silk cursor is moved *through* another 3D object, the user may estimate the object's depth in a number of ways, including estimating the distance traveled, timing, and kinesthesia.

4.2 Interactions Among Depth Cues: Modeling of 3D Performance

The manipulation of two sources of depth information in this experiment—occlusion and binocular disparity—brings to the fore two important theoretical questions: (1) when multiple sources of depth information are provided, how does the visual system judge actual depth information and (2) how does performance change accordingly? Our visual system could either select one of the multiple sources or integrate them to form a decision. Two classes of models have been applied to address this issue: additive models and multiplicative models [Bruno and Cutting 1988; Sollenberger 1993]. An additive model represents the fact that either depth cue can improve performance on its own, and when both sources of information are present simultaneously the resulting performance improvement is a simple summation of the benefit from the two sources individually. A multiplicative model describes the fact that the two sources of information can interact, causing a combined effect either greater or less than the additive effects. In their study of the combination of relative size, projection height, occlusion, and motion parallax, Bruno and Cutting [1988] concluded that additive models produced the best fit to their experimental data. In a series of experiments with motion parallax (kinetic depth) and binocular disparity, Sollenberger [1993] found some evidence for a multiplicative model with greater than additive effects for his path-tracing task.

In the present experiment, we found that task performance as measured by trial completion time and error rate were also compatible with a multiplicative model, but with less than additive effects. As shown in Figures 15 and 16, a strong interaction was found between display mode and cursor type for both trial completion time and error rate. That is, both stereo display alone (i.e., WireframeStereo) and partial occlusion alone (i.e., SilkMono) greatly improved performance relative to WireframeMono, but further improvements from SilkMono to SilkStereo (i.e., with both cues present) was marginal, suggesting the dominance of the partial-occlusion cue in this task.

For cases in which targets were missed, on the other hand, the pattern of error magnitudes (Figure 17) conformed with an additive model. No interaction was found between display mode and cursor type ($F_{1,11} = 0.0004$, $p = 0.97$).

5. FUTURE WORK

5.1 Future Research

Although this study has convincingly demonstrated some of the advantages of semitransparency applied to 3D human-computer interfaces, a number of issues related to semitransparency remain to be explored in the future. First, transparency is actually a continuous variable, ranging from total transparency to total occlusion. In the present experiment, the selected transparency value (determined by testing sample values during our pilot study) was compared against total transparency (the wireframe case). In practice, however, the optimal level of transparency may vary for different applications. Future work should compare all levels of transparency, including total occlusion (opaque).

Second, as mentioned before, partial occlusion may also be realized by drawing solid line segments (or cross hairs) on the cursor surface so that the cursor appears like a fishing net. This method also represents a continuum, along the dimension of line density. We expect that the partial occlusion provided by this “net” cursor approach will be inferior to the color interpolation method; however, a formal experiment would be worthwhile to carry out.

Finally, we used a dynamic target acquisition task to test the concept of semitransparency as a general interaction mechanism. Although independent of the theme of this article, an interesting issue related to the target acquisition task is the effect of relative size of the volume (or area) cursor versus the target. A separate study has been carried out in modeling such an effect through Fitts’ law and has been reported elsewhere [Kabbash and Buxton 1995].

5.2 Future Applications

In the background section, we have reviewed a number of existing user interfaces that make use of semitransparency. In the present section, we propose some other areas where we feel the interactive silk techniques could effectively be used. Our goal is to illustrate the potential diversity of such applications for future interface designs.

5.2.1 *3D Tool Glass.* The see-through interface widgets ("tool glass") developed by Bier et al. [1993] have proven to be quite advantageous in 2D interaction tasks [Kabbash et al. 1994]. The key concepts underlying "tool glass" include (1) making widgets (semi)transparent, so that the user can see through the widgets and superimpose them onto objects to be manipulated and (2) the user's two hands should be cooperatively involved in the interaction task (with the nondominant hand positioning the tool glass widgets and the dominant hand selecting items on the tool glass). In light of the results reported in the present article, extension of see-through interface widgets to 3D interactions is an obvious next step. According to this concept, a user would move a set of semitransparent tool glass widgets in 3D space with the nondominant hand while using the dominant hand in a coordinated fashion to complete the manipulation task.

5.2.2 *Virtual Reality and Augmented Reality.* Semitransparency is expected to see increasing usage in VR systems, where a large number of interactive widgets could be drawn with semitransparent surfaces. One example of this is the hand metaphor often used in VR applications as a representation of a user's own hand input image. Such a "cursor" can be either drawn in solid color or in wireframe. However, given the various manipulative functions of the hand representation (many of which involve occlusion of underlying objects), rendering the hand in semitransparency, as illustrated in Figure 21, is expected to be beneficial.

5.2.3 *Virtual Fixtures.* Virtual fixtures are a class of haptic and auditory aids proposed by Rosenberg [1993] as a means of improving teleoperation performance. We suggest that the notion of virtual fixtures can also be implemented using semitransparent graphical rendering, either combined with or independently of haptic and auditory aids. One of the examples of such a semitransparent fixture is a 2D plane that specifies the boundary beyond which movement of a robot arm becomes dangerous. Such a tool could off-load a human operator from the task of memorizing the locations of boundaries of similar real warning zones. Another advantage of a semitransparent virtual fixture is that the user could simultaneously maintain awareness of what lies on the other side of the warning plane.

5.2.4 *Telerobotic Control.* In order to off-load human operators from the task of continually controlling a telerobot in real time, some researchers have developed the technique of planning the robot's movements by means of a graphical/virtual robot model (e.g., Bejczy et al. [1990] and Zhai and Milgram [1991]). Such a "phantom robot" [Bejczy et al. 1990] is usually

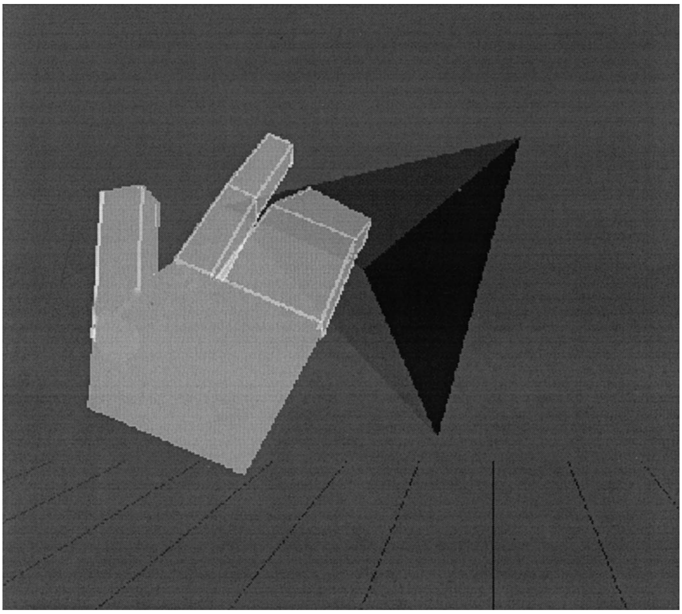

Fig. 21. A "silk magic hand" for VR applications.

drawn in a solid color or wireframe. A "silk phantom robot" (Figure 22) drawn in semitransparency could allow the operator to see objects behind the robot and to better visualize operations, particularly when the robot is in close proximity to obstacles and targets.

In conclusion, in this article we have proposed partial occlusion through semitransparency as a potentially powerful depth cue for computer interface applications, alongside such established 3D graphic techniques as perspective projection, stereoscopic displays, motion parallax, and viewpoint tracking. In an experimental investigation of the partial-occlusion cue, we have demonstrated its merits relative to the important stereoscopic cue in a 3D target acquisition task. For tasks in which *3D localization* is a critical component, semitransparency is expected to play a potentially very useful role in the future, not only in conventional computer graphic applications but also in such areas as telerobotic control and virtual reality.

ACKNOWLEDGMENTS

We would like to thank other members of the Input Research Group (IRG) (http://www.dgp.toronto.edu/IRG/) and the Ergonomics in Teleoperation and Control (ETC) Lab (http://vered.rose.toronto.edu) at the University of Toronto. We wish to thank George Fitzmaurice, Ken Hinckley, and Ferdie Poblete for their comments and assistance in conducting this research. We also wish to thank the anonymous referees of CHI'94 for their constructive comments on our earlier report of this work [Zhai et al. 1994], which

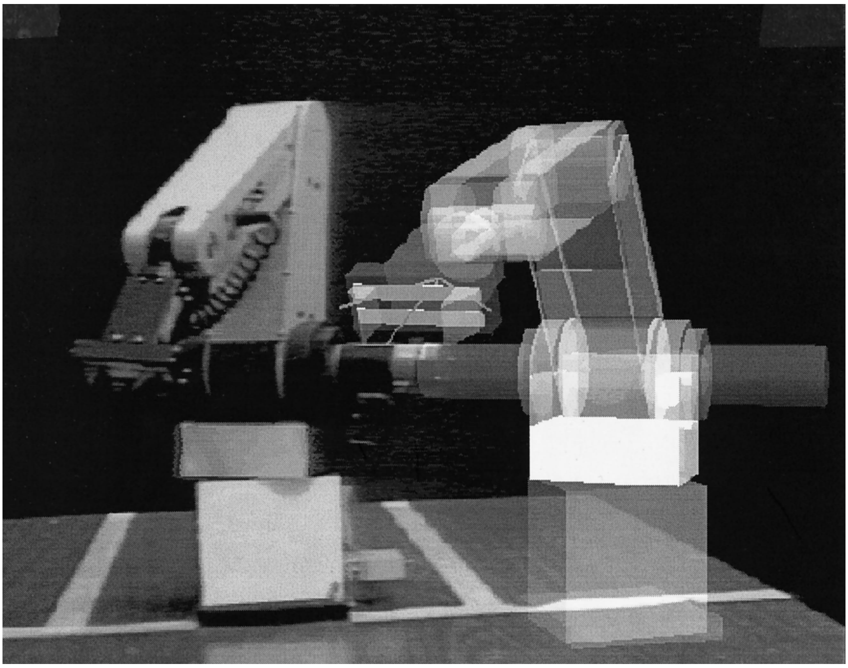

Fig. 22. The "silk phantom robot" for robot manipulation. Courtesy of Anu Rastogi.

motivated much of the reanalyses and discussions in this writing. The thorough and helpful reviews of the TOCHI referees have greatly improved the presentation of the article.

REFERENCES

ARDITI, A. 1986. Binocular vision. In *Handbook of Perception and Human Performance*, K. R. Boff, L. Kaufman, and J. P. Thomas, Eds. John Wiley and Sons. New York, 23–1, 23–41.

ARTHUR, K., BOOTH, K., AND WARE, C. 1993. Evaluating 3D task performance for fish tank virtual worlds. *ACM Trans. Inf. Syst. 11,* 3 (July), 239–265.

BEJCZY, A. K., KIM, W. S., AND VENEMA, S. C. 1990. The phantom robot: Predictive displays for teleoperation with time delay. In *Proceedings of IEEE International Conference on Robotics and Automation.* IEEE, New York, 546–551.

BIER, E. A., STONE, M. C., PIER, K., BUXTON, W., AND DEROSE, T. D. 1993. Tool glass and magic lenses: The see-through interface. In *Proceedings of SIGGRAPH 93.* ACM, New York.

BOCK, R. D. 1975. *Multivariate Statistical Methods in Behavioral Research.* McGraw-Hill, New York.

BRAUNSTEIN, M. L., ANDERSON, G. J., ROUSE, M. W., AND TITTLE, J. S. 1986. Recovering viewer-centered depth from disparity, occlusion and velocity gradients. *Perception Psychophys. 40,* 216–224.

BROOKS, F. P. J. 1988. Grasping reality through illusion—Interactive graphics serving science. In *Proceedings of CHI'88: ACM Conference on Human Factors in Computing Systems.* ACM, New York.

BRUNO, N. AND CUTTING, J. E. 1988. Minimodularity and the perception of layout. *J. Exper. Psychol. Gen. 117,* 161–170.

CARD, S., ROBERTSON, G., AND MACKINLAY, J. 1991. The information visualizer. In *Proceedings of CHI'91: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 181–194.

CARD, S. K., PIROLLI, P., AND MACKINLAY, J. D. 1994. The cost-of-knowledge characteristic function: Display evaluation for direct-walk dynamic information visualizations. In *Proceedings of CHI'94: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 238–244.

CHEN, M., MOUNTFORD, S. J., AND SELLEN, A. 1988. A study in interactive 3-D rotation using 2-D control devices. In *Proceedings of ACM SIGGRAPH'88.* ACM, New York, 22.

CROW, F. AND PIZER, S. M., Eds. 1986. *Proceedings of the 1986 Workshop on Interactive 3D Graphics.* ACM, New York.

ELLIS, S. R., KAISER, M. K., AND GRUNWALD, A. J., Ed. 1991. *Pictorial Communication in Virtual and Real Environments.* Taylor and Francis, London.

FOLEY, J. D., VAN DAM, A., FEINER, S. K., AND HUGHES, J. F. 1990. *Computer Graphics Principles and Practice.* Addison-Wesley, Reading, Mass.

HABER, R. N. AND HERSHENSON, M. 1973. *The Psychology of Visual Perception.* Holt, Rinehart and Winston, New York.

HARRISON, B. L., ISHII, H., VICENTE, K. J., AND BUXTON, W. 1995. Evaluation of a display design space: Transparent layered user interfaces. In *Proceedings of CHI'95: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 317–324.

HARRISON, B. L., ZHAI, S., VICENTE, K. J., AND BUXTON, W. 1994. Semi-transparent "silk" user interface objects: Supporting focused and divided attention. CEL Tech. Rep., Dept. of Industrial Engineering, Univ. of Toronto, Ontario, Canada.

HERNDON, K. P., ZELAZNIK, R. C., ROBBINS, D. C., CONNER, D. B., SNIBBE, S. S., AND VAN DAM, A. 1992. Interactive shadows. In *Proceedings of the ACM Symposium on User Interface Software and Technology.* ACM, New York, 1–6.

HINCKLEY, K., PAUSCH, R., GOBLE, J. C., AND KASSELL, N. F. 1994a. A survey of design issues in spatial input. In *Proceedings of the ACM Conference on User Interface Software and Technology.* ACM, New York.

HINCKLEY, K., PAUSCH, R., GOBLE, J. C., AND KASSELL, N. F. 1994b. Passive real-world interface props for neurosurgical visualization. In *Proceedings of CHI'94: ACM Conference on Human Factors in Computing Systems.* ACM, New York.

HOWELL, D. C. 1992. *Statistical Methods for Psychology.* 3rd ed. PWS-Kent, Boston, Mass.

JACOB, R. J., K., SIBERT, L. E., MCFARLANE, D. C., AND MULLEN, M. P. 1994. Integrality and separability of input devices. *ACM Trans. Comput. Human Interac. 1,* 1 (Mar), 3–26.

KABBASH, P. AND BUXTON, W. 1995. The "Prince" technique: Fitts' law and selection using area cursor. In *Proceedings of CHI'95: ACM Conference on Human Factors in Computing Systems.* ACM, New York.

KABBASH, P., BUXTON, W., AND SELLEN, A. 1994. Two-handed input in a compound task. In *Proceedings of CHI'94: ACM Conference on Human Factors in Computing System.* ACM, New York, 417–423.

KAUFMAN, L. 1974. *Sight and Mind—An Introduction to Visual Perception.* Oxford University Press, London.

LIANG, J. AND GREEN, M. 1994. JDCAD: A highly interactive 3D modeling system. *Comput. Graph. 18,* 4, 499–506.

MCALLISTER, D. F., Ed. 1993. *Stereo Computer Graphics and Other True 3D Technologies.* Princeton University Press, Princeton, N.J.

OLDFIELD, R. C. 1971. The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia 9,* 97–113.

OVERBEEKE, C. J. AND STRATMANN, M. H. 1988. Space through movement. Ph.D. thesis, Delft Univ. of Technology, Delft, Netherlands.

POULTON, E. C. 1974. *Tracking Skill and Manual Control.* Academic Press, New York.

ROBERTSON, G. G., MACKINLAY, J. D., AND CARD, S. K. 1991. Cone trees: Animated 3D visualizations of hierarchical information. In *Proceedings of CHI'91: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 189–194.

ROSENBERG, L. B. 1993. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *Proceedings of IEEE Virtual Reality Annual International Symposium (VRAIS'93).* IEEE, New York, 76–82.

SCHRIEVER, W. 1925. Experimentelle studien über das stereoskopische sehen. *Zeitschrift für Pscyhologie 96,* 113–170.

SHERIDAN, T. B. 1992. *Telerobotics, Automation, and Human Supervisory Control.* The MIT Press, Cambridge, Mass.

SMETS, G. J. F. 1992. Designing for telepresence: The interdependence of movement and visual perception implemented. In *Proceedings of 5th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-Machine Systems*, Pergamon Press.

SOLLENBERGER, R. L. 1993. *Combining depth information: Theory and implications for design of 3D displays.* Ph.D. thesis, Dept. of Psychology, Univ. of Toronto, Ontario, Canada.

SOLLENBERGER, R. L. AND MILGRAM, P. 1993. Effects of stereoscopic and rotational displays in a three-dimensional path-tracing task. *Human Factors 35,* 3, 483–499.

VENOLIA, D. 1993. Facile 3D direct manipulation. In *Proceedings of INTERCHI'93: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 31–36.

WARE, C. 1990. Using hand position for virtual object placement. *Vis. Comput. 6,* 245–253.

WARE, C. AND ARTHUR, K. 1993. Fish tank virtual reality. In *Proceedings of INTERCHI'93: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 37–42.

WICKENS, C. D., TODD, S., AND SEIDLER, K. 1989. Three-dimensional displays: Perception, implementation and applications. CSERIAC Tech. Rep. 89-001, Wright Patterson Air Force Base, Ohio.

YEH, Y. Y. 1993. Visual and perceptual issues in stereoscopic display. In *Stereo Computer Graphics,* D. F. McAllister, Eds. Princeton University Press, Princeton, N.J., 50–70.

YEH, Y. Y. AND SILVERSTEIN, L. D. 1992. Spatial judgments with monoscopic and stereoscopic presentation of perspective displays. *Human Factors 34,* 5, 583–600.

ZELTZER, D. 1992. Autonomy, interaction, and presence. *Presence 1,* 1, 127–132.

ZHAI, S. 1995. Human performance in six degrees of freedom input control. Ph.D. thesis, Univ. of Toronto, Ontario. Available via http://vered.rose.toronto.edu/people/shumin.html.

ZHAI, S. AND MILGRAM, P. 1991. A telerobotic virtual control system. In *Proceedings of SPIE: Cooperative Intellegent Robotics in Space II.* Vol. 1612. SPIE, Boston, Mass.

ZHAI, S. AND MILGRAM, P. 1993. Human performance evaluation of manipulation schemes in virtual environments. In *Proceedings of VRAIS'94: The 1st IEEE Virtual Reality Annual International Symposium.* IEEE, New York.

ZHAI, S. AND MILGRAM, P. 1994. The "silk cursor": Investigating transparency for 3D target acquisition. In *Proceedings of CHI'94: ACM Conference on Human Factors in Computing Systems.* ACM, New York, 459–464.

Received December 1994; accepted January 1996