Light-Emitting Cursors for 3D Navigation

Avi C. Naiman
Department of Psychology
New York University

Cheuk-Mun Ada Yuen
National Computer Systems

Contact Author:

Avi C. Naiman
Department of Psychology
New York University
6 Washington Place, Eighth Floor
New York, NY 10003
(212) 998-3880
FAX: (212) 995-4349
avi@cns.nyu.edu

Abstract

The lack of relative depth information in 3D environments displayed on 2D computer screens hinders 3D navigation. To assist users in this task, we have developed a light-emitting cursor (LEC) – a 3D cursor with a light source attached to its focus point. The light from a LEC provides relative depth information between an object and the LEC, as well as amongst objects in the scene. Results from a traversal experiment demonstrate the effectiveness of a LEC to reduce the time and traversal distance required to navigate in three dimensions.

CR Categories: I.3.6 [**Computer Graphics**]: Methodology and Techniques – *interaction techniques.* H.5.2 [**Information Interfaces and Presentation**]: User Interfaces – *input devices and strategies, interaction styles, screen design*; H.1.2 [**Models and Principles**]: User/Machine Systems – *human factors*;

General Terms: 3D navigation, user interface design.

Additional Keywords and Phrases: light-emitting cursors.

Introduction

The past few years have brought an explosion in the availability of 3D applications, whether for design, visualization, education or entertainment. Although, human-computer interface designers have considerable experience mapping 2D input and output devices with 2D environments and interaction techniques (e.g., [Fitts 54, MacKenzie *et al.* 91, MacKenzie and Buxton 93]), with the introduction of the third dimension, either new hardware – e.g., virtual-reality helmets and Spaceballs™ – must be acquired and mastered or existing 2D devices and techniques must be adapted (see [Chen *et al.* 88, Evans *et al.* 81, Herndon *et al.* 94, Mackinlay *et al.* 90, Nielson and Olsen 86]). Due to the costs of novel hardware, the vast majority of users will, for some time to come, have to depend on 2D displays, even for 3D applications.

Two common tasks in such applications are 3D *navigation* – moving the point of focus, often represented by a visual *cursor* – and 3D *selection* – choosing objects in a scene. Inherent in these tasks is the need for a user to determine the 3D locations of the objects, including the cursor. The primary mechanism for conveying 3D location information is the visual system's ability to perceive depth. While human depth perception is far from completely understood, depth cues include: perspective, object size, retinal image size, relative brightness, diffuse and specular shading, shadowing, occlusion, motion parallax, object rotation, binocular disparity, accommodation, and convergence (see [Braunstein *et al.* 86, Kaufman 74, Landy *et al.* 95]).

Previous studies have reported on visual and navigational aspects of virtual-reality systems (e.g., [Deering 92, Pausch *et al.* 95]). However, the prevalence of such systems is still quite low. Consequently, we have focused our efforts on aiding the user's perception of depth information when no stereoscopic depth cues are available; i.e., a scene is presented on a typical 2D display surface and viewed without any special glasses. We are particularly interested in software techniques that provide relative depth information using standard hardware.

Any techniques that can make 3D navigation and selection easier, without overly burdening rendering systems, can prove quite beneficial. Along these lines, recent research has resulted in a number of 3D cursors to improve 3D *selection*:

- The Beam Cursor uses beams of light emanating from the cursor to allow quick selection of distant objects by orienting the cursor in the direction of interest, without needing to reposition the cursor [Yoshimura *et al.* 94]. Relative depth information between an object and the cursor is determined from the cursor's orientation and the beam-object intersection. However, no relative depth information amongst objects in the scene is available. Furthermore, the usefulness of the Beam Cursor declines as the number of objects increases.
- The Spotlight Cursor is similar to the Beam Cursor, but uses a cone of light rather than a focused beam [Liang and Green 94]. This allows for easier selection of distant objects with less precise orientation than required by the Beam Cursor. However, the Spotlight Cursor also suffers from the lack of relative depth information amongst objects in the scene and its declining usefulness with an increase in the number of objects.
- The Apple Cursor takes the form of a cone or jack, depending on whether it is inside an object [Venolia 93]. Relative depth information between an object and the cursor, as well as amongst objects, is determined from shadows cast on a platform below all the objects in the scene. However, when shadows are hidden from the user's point of view by other objects, the relative depth information is lost. Furthermore, this design necessitates a platform placed beneath all objects.
- The Silk Cursor uses a semi-transparent, box-shaped cursor and occlusion by the cursor's front and back surfaces to provide relative depth information between an object and the cursor [Zhai *et al.* 96]. However, only two levels of semi-transparent occlusion are available and there is no relative depth information for any set of objects that are all either inside or outside of the cursor volume.

All of these cursors were designed primarily to aid in 3D selection, rather than 3D navigation. As a result, they are limited in the 3D information that is universally available. For the purposes of 3D navigation, relative depth information between the cursor and an object of interest – as well as amongst all objects – should be persistent and obvious.

Here, we introduce a novel cursor we have developed specifically to aid in 3D *navigation*. To test its usefulness, we carried out a 3D navigation experiment, measuring a variety of performance measures – such as completion time and distance traveled – and comparing performance both with and without size and occlusion depth cues. In addition to the specific improvements demonstrated with this cursor, the experimental paradigm introduced below also serves as a general template for investigating 3D user-interface components.

Light-Emitting Cursors

Our approach to improving the user's depth perception of a 3D scene is to turn the cursor into a local, omnidirectional light source; we refer to such a cursor as a *light-emitting cursor* (LEC). As the cursor is moved in the 3D environment, the pattern of light it casts on other objects changes *continuously*, providing visual feedback as to the 3D position of the cursor, as well as to the relative 3D positions of the other objects in the scene.

We designed our LEC to be a cone with a point light source at its tip, radiating light equally in all directions. The tip of the cone also serves as the 'position' of the cursor for the purpose of capturing objects. Although the conical shape provides an intuitive direction vector for navigation (but not orientation information around the direction vector), we are concerned here only with the position of the cursor and permit navigation in any direction.† While we could have designed a cursor without a directional component, we

† In fact, subjects in the experiment to be described below often preliminarily oriented the cursor in the direction of planned navigation. However, as this was true whether or not the cursor emitted light, there is no biased effect on the results reported.

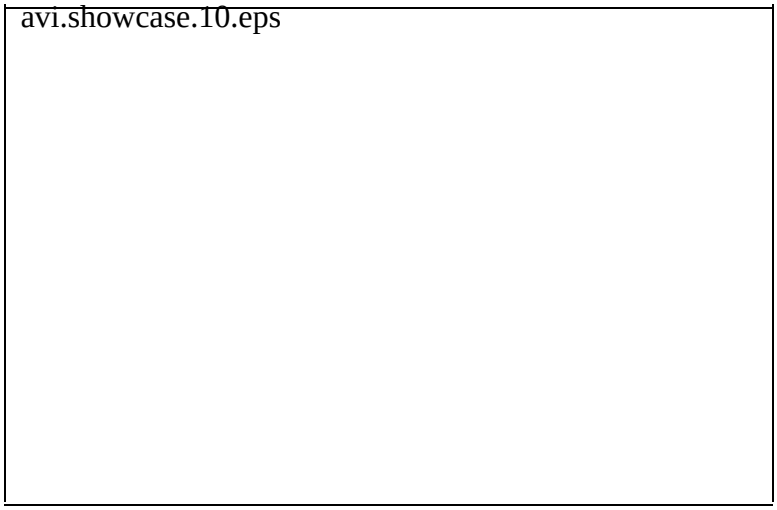

Figure 1: Relative Depth Between Cursor and Sphere (left: LEC; right: NLEC).

chose such a shape to allow us to expand future investigations to include directional considerations without having to confound the experimental design with a new cursor shape.

Hypothesized Advantages

Fig. 1 demonstrates the utility of a LEC in determining whether the cursor is closer to or further from the viewer than some other object. On the top-left, a LEC is further away from the viewer than the sphere, resulting in a small highlight along the right edge of the sphere, in addition to the highlight on top of the sphere, which comes from a static, overhead light source. On the bottom-left, the LEC is closer than the sphere, resulting in a large highlight on the front of the sphere, in addition to the same highlight on top of the sphere. On the right side of the figure are identical placements of the cursor and sphere, but this time the cursor is non-light-emitting (NLEC); note the lack of visual cues with respect to relative depth.

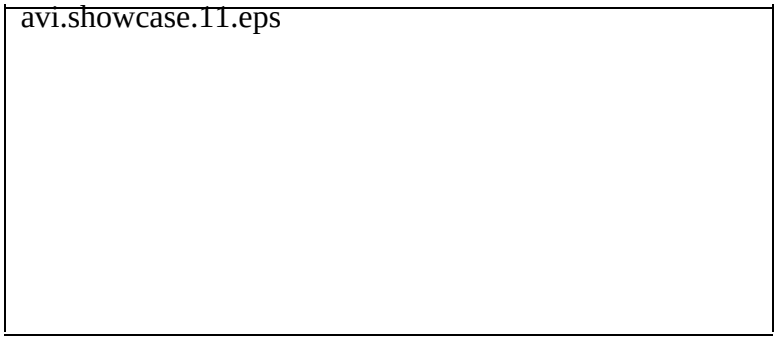

Figure 2: Relative Depth Amongst Spheres (left: LEC; right: NLEC).

Similarly, the reflections from a LEC can provide relative depth information amongst a collection of objects in the scene (Fig. 2), which might also aid in navigation planning. On the left, the LEC gives a clear perception of the depth ordering of the three spheres, whereas relative depth is difficult to determine on the right, where a NLEC was used. Note that, if it is known *a priori* that all three spheres are the same *physical* size, then the *projected* sizes can serve as a reliable depth cue, even with the NLEC; this potential confounding depth cue will be addressed below. Of course, for objects that vary in physical size, projected size alone is not a dependable depth cue.

In addition to depth perception, the use of a LEC can also aid in distance estimation, as seen in Fig. 3. In the series of images on the left, a LEC is moved closer to the sphere, from top to bottom. As the LEC approaches the sphere, the highlight due to the LEC becomes smaller and interacts less with the highlight from the overhead light. On the right, the same approach using a NLEC shows no such distance cue.

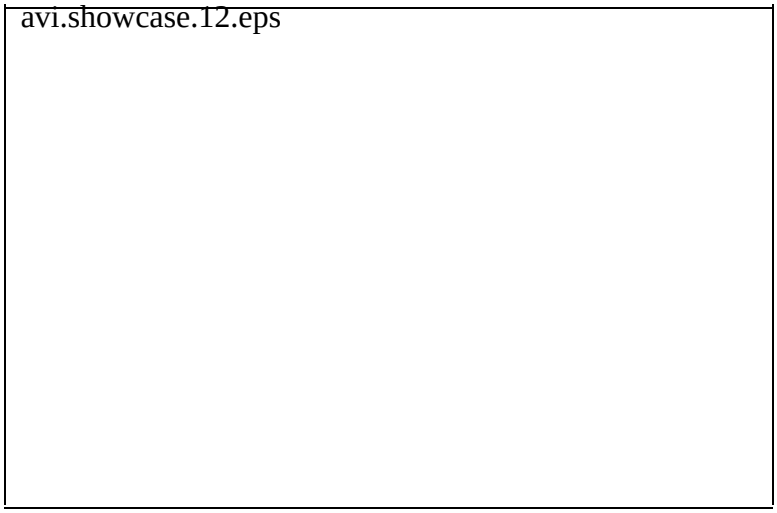

Figure 3: Relative Distance Between Cursor and Sphere (left: LEC; right: NLEC).

Navigation Experiment

To evaluate the effectiveness of having a cursor emit light, we had 4 subjects (one of whom was also an experimenter) carry out a navigation task with both a LEC and a NLEC and analyzed a number of performance measures, such as completion time and distance traveled (see below). All subjects had extensive prior experience with 2D computer systems, but none had any prior experience with 3D systems. All subjects had corrected-to-normal vision, were familiar with the goals of the experiment, and had previous experience participating in vision and/or interaction experiments.

The experiment was carried out on a SGI Indigo^2 Extreme workstation with a Mitsubishi color monitor (model HL7965KW-SG); the input device was a 6 degree-of-freedom Spaceball™.† Subjects were free to position their chairs and the Spaceball™ so as to make

† While most users will have available to them only 2D input devices, the physical and cognitive mappings from such devices to a 3D space are themselves complex issues and are unrelated to our interest here in the LEC. Consequently, we chose a 3D input device available to us and ensured that subjects had sufficient training in its use (see below).

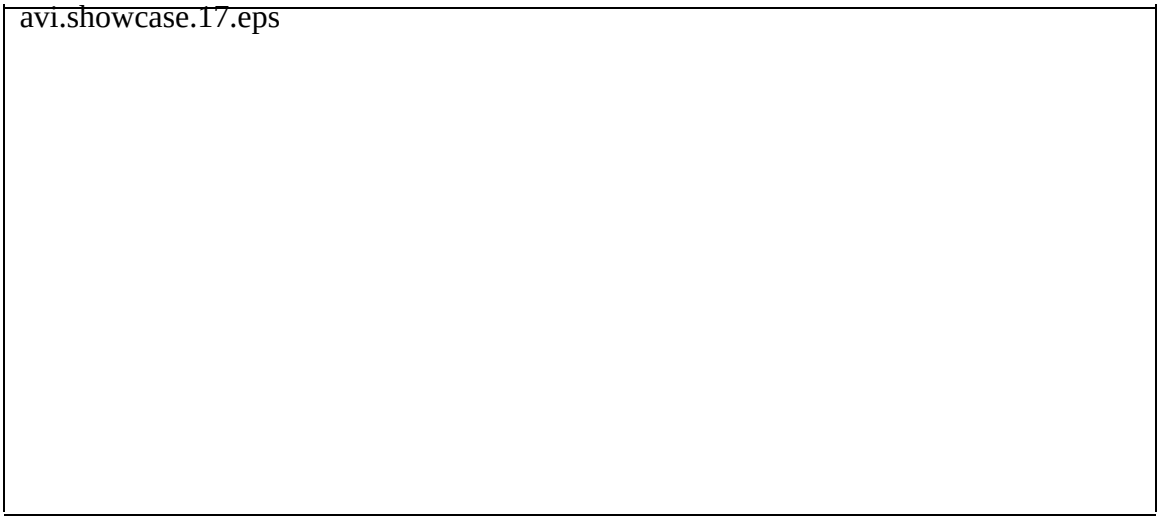

Figure 4: Sample Trial in Navigation Experiment (left: LEC; right: NLEC).

themselves comfortable. Before each series of trials, the computer screen was calibrated so that a white screen produced 80 cd/m^2 at its center and luminance was linear in DAC setting (i.e., gamma-corrected). Other than the light provided by the computer screen, the experimental environment was kept dark.

Fig. 4 shows a sample static image from a trial with the LEC (left) and NLEC (right). In each trial, 10 gray spheres and the cursor (also gray) were placed randomly in a (black) 3D design space, constrained so that they were all completely visible in the scene's 2D projection onto the computer screen. Spheres were deliberately chosen because their lighted surfaces provide optimal depth information and, hence, the greatest likelihood of demonstrating an effect for LECs; consequently, any benefits realized using spheres serve as upper-bound improvements. A static light source was placed above the center of the design space (but outside of the projected image); this served as the only light source when using a NLEC and as a secondary light source when using a LEC.

A hardware-implemented, industry-standard lighting model was used (for further explication of the model, see [Foley *et al.* 90, McLendon 92]):

$$I = I_a k_a + \sum_{1 \le i \le m} I_{a_i} k_a + f_{att_i} I_{p_i} \left[k_{d_i} (\overrightarrow{N} \cdot \overrightarrow{L_i}) + k_{s_i} (\overrightarrow{R_i} \cdot \overrightarrow{V})^n \right],$$

where f_{att_i}, the attenuation factor for point light source (PLS) *i*, is defined as:

$$\min\left(\frac{1}{c_1 + c_2 d_{L_i} + c_3 {d_{L_i}}^2}, 1 \right).$$

The meanings and settings for the various lighting model symbols are listed in Table 1. The settings used in the experiment were chosen by trial and error to ensure good visibility of the spheres regardless of location in the design space as well as reasonable attenuation of the LEC's light as the cursor was moved around within the space. The shading model used was Gouraud [Foley *et al.* 90] and there were no secondary reflections. These hardware-implemented lighting and shading models do not include shadowing. That is, the light computed at any surface does not take into consideration any opaque objects that intervene between the surface and the light source. This will be further addressed below.

The subject's task was to 'touch' each of the 10 spheres as quickly as possible, by positioning the tip of the cursor anywhere within the volume of each sphere. Although the task involves 3D *selection* of a set of objects, note that the experiment is designed specifically to elicit information about the 3D *navigation* required to reach all of the objects. The subject manipulated the position of the cursor through the use of a Spaceball™; the subject was also free to manipulate the cursor's 3D orientation, but was not instructed in this regard. To ensure that spheres were not 'touched' merely by chance, subjects were required to press a button on the Spaceball™ when they believed that the tip of the cursor was inside a sphere. If correct, the sphere was marked as touched by turning its color red, but ensuring that the luminance remained unchanged. A trial terminated when all 10 spheres had been touched.

Symbol	Meaning	Settings
I	resulting intensity	computed
I_a	intensity of global ambient component	0.0
k_a	ambient-reflection coefficient	0.3
m	number of PLSs	NLEC: 1; LEC: 2
I_{a_i}	intensity of ambient component for PLS i	NLEC: 0; LEC: 0.0; Overhead Light: 0.3
I_{p_i}	intensity of PLS i	NLEC: 0; LEC: 0.3; Overhead Light: 0.9
k_{d_i}	diffuse-reflection coefficient for PLS i	1.0
$\vec{N}$	surface normal	computed
$\vec{L_i}$	vector to PLS i	computed
k_{s_i}	specular-reflection coefficient for PLS i	0.5
$\vec{R_i}$	reflection vector for PLS i	computed
$\vec{V}$	viewpoint vector	computed
n	specular-reflection exponent	100
d_{L_i}	distance to PLS i	computed
c_1, c_2, c_3	attenuation coefficients	1, 0.0185, 0

Table 1: Settings for Lighting Model Parameters.

While the depth cues provided by a LEC undoubtedly combine with any other available depth cues in complex ways [Landy *et al.* 95], we were interested here in investigating the maximal utility of a LEC. Consequently, we isolated the depth cue resulting from use of a LEC by removing others:

- *Object motion.* Throughout a trial, only the cursor moved; neither the subject's viewpoint nor the positions of the spheres varied.
- *Projection Size.* If all of the spheres would have had the same *physical* size, their projections onto the 2D display screen would have varied in size. Consequently,

once a (random) position for each sphere was chosen, its physical size was manipulated to provide a constant-sized projection onto the screen. However, it would be unnatural for the cursor's projected size to remain constant as it moved about in 3D, so its physical size was fixed and its projected size varied with distance from the projection plane.

- *Occlusion*. The randomized sphere positions were constrained to preclude occlusion. Occlusion between the cursor and the spheres was, of course, allowed.

A set of 50 different trials was created using the LEC and the same set of 50 trials was repeated using the NLEC. As we will discuss below, 6 further sets of 50 trials (3 for LEC and 3 for NLEC) were used in the experiment, resulting in a total of 400 trials per subject. For each subject, the presentation order of the 400 trials was randomized and the trials were broken into 8 blocks of 50. Each block was preceded by 10 practice trials and took about 45 minutes to complete. The experiment was completed by each subject within a 2-week period.

Training

Although all subjects had extensive experience with 2D computer systems, the novelty of the input device, interaction techniques for navigation, experimental environment, depth cues and task necessitated a fair amount of familiarization. Training consisted of 10 20-trial sessions (one per day). The first 6 sessions used the experimental condition described above; the final 4 sessions used conditions to be described below. The final 10 trials of each session were recorded and examined for learning trends.

An analysis of the means and standard errors of a variety of performance measures (see below) suggested that, after 10 days of training, the effect of a subject's proficiency in using the Spaceball™ no longer had a significant impact on performance. For example, Fig. 5 plots the mean completion times and ±1 standard errors for one subject. For more details, see [Yuen 96].

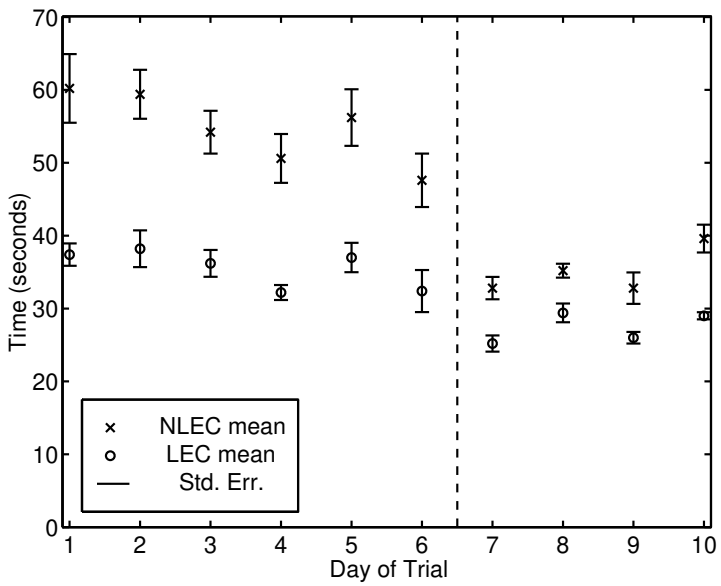

Figure 5: Completion Times as a function of Training.

Performance Measures

We will compare performance when using a LEC with performance when using a NLEC for the following measurement variables (see Fig. 6):

- *Completion Time*: the time taken to complete a trial.
- *Total Distance Traveled*: the distance traveled by the cursor from its initial position until the trial completes; the dashed lines in the figure.
- *Total Theoretical Distance*: the straight-line distance that could, in theory, have been achieved, for the *particular order* in which the subject touched the spheres, when considering minimum distances only from each sphere to the next (i.e., this is *not* a measure of the optimal traversal path, which might require a tangential touch or even use a different touch order); the solid, black lines in the figure. This measure tells us about the subject's ability to *plan* an optimal traversal path: a shorter Total

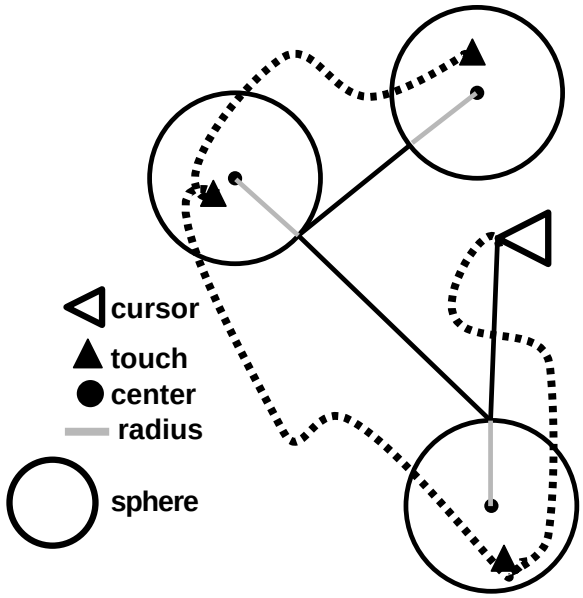

Figure 6: Sample Traversal Path.

Theoretical Distance implies the subject had better perception of the relative depths of the spheres.

- *Ratio of Total Distance Traveled to Total Theoretical Distance*: this measure tells us how much a subject wandered from the straight-line path, which may also have consisted of a change of target along the way from one sphere to another.
- *Ratio of* Final *Distance Traveled to* Final *Theoretical Distance*: like the previous measure, this one tells us how much a subject wandered from the straight-line path between the final two spheres. However, because there is only the final sphere available as a target after the penultimate sphere, there cannot have been a change of target along the way. Consequently, this is a more ‘pure’ measure of the subject’s ability to travel in straight lines.

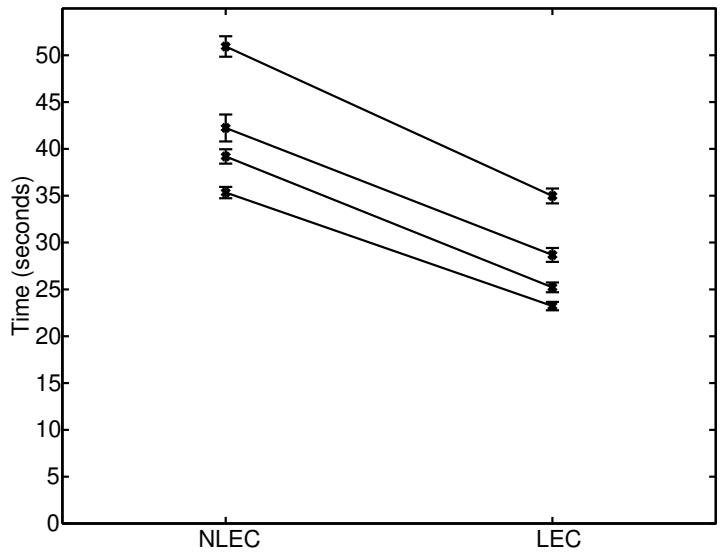

Figure 7: Completion Times for Each of the 4 Subjects.

Results

Fig. 7 plots the mean completion times and ±1 standard errors for all 4 subjects when using the NLEC and LEC. Even without statistical analysis, it is quite apparent from the large separations between corresponding means and the small standard errors that there is a significant main effect of the cursor serving as a light source. Furthermore, although there are *absolute* differences in performance between subjects, the *relative* improvements (i.e., the percentage reductions in completion time) when using the LEC compared with the NLEC are quite similar across subjects: 31.33%, 32.10%, 35.66% and 34.30%, from top to bottom.

To determine if there was a main effect of the cursor serving as a light source, we carried out an analysis of variance (ANOVA) for each of the performance measures. The results for almost all subjects and performance measures exhibited a reliable dependence

Performance Measure	Grand Mean Improvement	95% Conf. Interval
Completion Time	33.35%	3.16%
Total Distance Traveled	34.45%	6.06%
Total Theoretical Distance	24.08%	4.37%
Ratio of Total Dist. Traveled to Total Theoretical Dist.	14.53%	9.00%
Ratio of Final Dist. Traveled to Final Theoretical Dist.	21.76%	7.58%

Table 2: Improvements in Performance Measures for LEC Compared with NLEC.

on cursor type ($F_{1,98}$, $p < 0.03$). The only exception was for 1 subject for the Ratio of Total Distance Traveled to Total Theoretical Distance measure. Table 2 lists the grand mean improvements across subjects for the various performance measures when using a LEC as opposed to a NLEC, as well as the 95% confidence intervals of the grand means; these are our best estimates of an untested user's performance.

Strategies

In follow-up interviews, subjects reported the navigation strategies they used for choosing a target sphere, approaching it, and touching it. Recall that subjects were requested to touch all spheres as quickly as possible. In the NLEC trials, path planning was impossible, so subjects chose the 'closest' sphere based solely on the x- and y-positions. Touching the sphere was accomplished by occlusion between the cursor and the sphere, followed by depth adjustment until the cursor touched the sphere.

In the LEC trials, path planning was possible: subjects used the lighted areas of spheres to guide their choice of the next target. However, the change in the highlights as the cursor moved occasionally resulted in a change of target. A straight-line approach to

a target was guided by the lighted area on the sphere, while fine movement was guided by the shrinking of the lighted area.

The ability to plan a path accounts for improvements in Completion Time, Total Distance Traveled, and Total Theoretical Distance. The ability to approach a target in a straight line accounts for improvements in Completion Time, Total Distance Traveled, Ratio of Total Distance Traveled to Total Theoretical Distance, and Ratio of Final Distance Traveled to Final Theoretical Distance.

Interaction With Other Depth Cues

As mentioned above, to learn as much as possible about the impact the LEC might have, we deliberately removed depth cues related to occlusion amongst spheres as well as to the spheres' projected sizes. To test the interactions between these cues and the effect of the LEC, each subject ran 3 additional sets of 100 trials (50 with the LEC and 50 with the NLEC) for each of the following 3 conditions: II) the projected sizes were allowed to vary (i.e., by maintaining constant physical sphere size); III) occlusion amongst spheres was permitted; and IV) both (Fig. 8). As mentioned above, these 300 trials were actually intermingled with the 100 trials of the original condition, forming the overall 400 trials presented to each subject. Table 3 lists the grand mean improvements across subjects for the various performance measures when using a LEC as opposed to a NLEC, for each of the three additional experimental conditions.

An ANOVA again found a reliable dependence on cursor type for almost all subjects, conditions and performance measures, with the following notable exception: when the projected sphere size is allowed to vary, there is no reliable dependence on cursor type for the Total Theoretical Distance measure. This was true for all 4 subjects and regardless of whether or not occlusion was allowed. Apparently, the depth cue from varying projected sphere size allows as good path planning as the LEC, and combining the two cues yields no further improvement. However, note that, for all other performance measures,

avi.showcase.20.eps

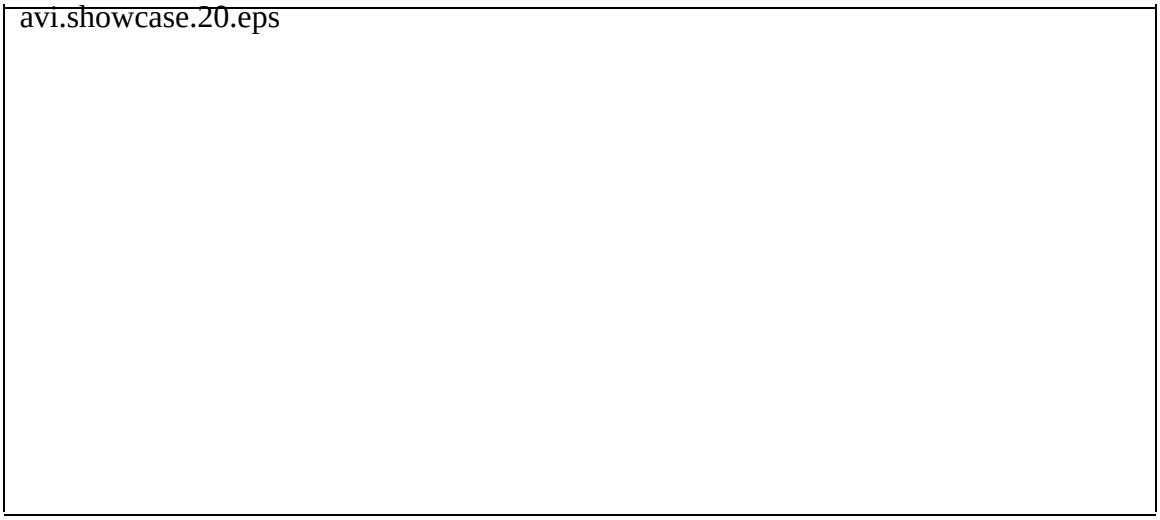

Figure 8: Sample Trial When Projected Size Varies and Occlusion is Allowed (left: LEC; right: NLEC).

there was a significant improvement even when projected size varied. That is, even though path planning showed no significant improvement, both completion time and distance traveled still exhibited significant and substantial improvements (though somewhat reduced from the improvements achieved in the original condition). Hence, use of a LEC provides a depth cue and associated performance improvements that are not otherwise available.

Conclusions

To aid in 3D navigation using a 2D display device, we have designed a cursor that emits light. The experiment described here clearly demonstrates the utility of the LEC as a 3D tool for judging relative position and for homing in on objects, even when combined with other depth cues such as projected size and occlusion. It is particularly noteworthy that improvements were obtained for a variety of performance measures. In the absence

Performance Measure	Grand Mean Improvement		
	Cond. II	Cond. III	Cond. IV
Completion Time	26.58%	35.65%	28.03%
Total Distance Traveled	21.63%	37.86%	23.64%
Total Theoretical Distance	0.93%	27.51%	1.99%
Ratio of Total Dist. Traveled to Total Theoretical Dist.	21.08%	15.63%	22.49%
Ratio of Final Dist. Traveled to Final Theoretical Dist.	23.94%	19.52%	30.86%

Table 3: Improvements in Performance Measures for LEC Compared with NLEC.
Condition II) Projected Size Varies. Condition III) Occlusion is Allowed. Condition IV) Both.

of other depth cues, one can expect approximately 26% shorter traversal paths, 36% shorter distances traveled, and 35% shorter completion times. When a projected-size depth cue is also provided, improvements drop to 23% shorter distances traveled and 27% shorter completion times, and traversal paths are similar with or without light-emitting cursors.† We expect that these 3D navigation benefits would be just as noticeable by users employing 2D input devices, although the navigation process would be more cumbersome, both physically and cognitively.

Recall that the four alternative cursors mentioned in the Introduction were designed primarily to aid in 3D selection; they offer little direct aid in 3D navigation. The LEC, on the other hand, was expressly designed to aid in 3D navigation: it provides omnipresent, continuously-varying depth information, both between an object and the cursor, as well as amongst other objects in the scene.

† Note, however, that, in order for a projected-size depth cue to have such an impact, users must already be aware of the objects' true physical sizes.

Future Work

Potential extensions to the basic LEC design presented here include user control of the LEC's intensity. For example, to guarantee visible effects in a brightly-illuminated scene, the user could increase the relative strength of the LEC. Conversely, if the user wants to examine the scene without interference from the LEC, a simple switch could allow for turning it off.

Of particular interest would be to combine the LEC's navigation aid with the selection assistance provided by the alternative 3D cursors mentioned earlier. For example, an additional (directional) spot light (perhaps of a different color) could be incorporated to assist in quick object selection. Such an extension would benefit from investigation of the cursor's orientation controls as well.

Recall that the hardware-implemented lighting and shading models did not include shadowing, resulting in potentially conflicting depth cues. Yet, despite this, we obtained significant improvements when using the LEC. We expect that correct treatment of shadows (see, for example, [Hudson 92]) would only increase the relative improvement provided by a LEC, by removing conflicting depth cues. This effect remains to be investigated, as do interactions with other depth cues such as object motion.

The experimental design used here isolated the effect of a LEC in a simple 3D navigation task. Furthermore, highly-reflecting spheres were chosen for the high degree of depth information they impart. The generalizability of these results to diffusely-reflecting and/or non-spherical objects requires further investigation. Similarly, the utility of such cursor designs must be tested in realistic applications, such as CAD/CAM, animation design and information visualization. More research is also needed to map out the effects of various lighting-model parameters as well as the use of multiple light sources. A host of other factors worthy of exploration include: the use of different cursor and object shapes, surface textures, alternative shading models, transparency, and wireframe drawings.

The novel idea of improving 3D navigation by having the cursor emit light opens up a host of research avenues. While it is unclear at this point which depth cues are best or how best to combine them, further work in this area should yield significant improvements to the problems of navigation and selection in 3D spaces. We hope that the methodologies presented here can serve as a template for investigating performance aspects of 3D user-interface components.

Acknowledgements

This research was supported under RGC grant CERG HKUST600/94E (Hong Kong). We thank: Tom Davis for sending us an old GL Manual; Dan Venolia for providing images and explanations of his cursor designs; Shumin Zhai for providing images of the Silk Cursor, helpful feedback on the experimental design and assistance with the Spaceball™; David Lam, Nelson Mak and Annette Werner for participating in the experiment; Michael S. Landy for helpful comments; and the Department of Psychology at the New York University for providing facilities with which to prepare this manuscript.

References

Braunstein *et al.* 86

M. L. Braunstein, G. J. Anderson, M. W. Rouse, and J. S. Tittle, ‘‘Recovering Viewer-Centered Depth from Disparity, Occlusion and Velocity Gradients,’’ *Perception & Psychophysics*, Volume 40, 1986, pp. 216-224.

Chen *et al.* 88

Michael Chen, S. Joy Mountford, and Abigail Sellen, ‘‘A Study in Interactive 3-D Rotation Using 2-D Control Devices,’’ *Computer Graphics*, Volume 22, Number 4, August 1988, pp. 121-128. SIGGRAPH 1988 Proceedings.

Deering 92

Michael Deering, ‘‘High Resolution Virtual Reality,’’ *Computer Graphics*, Volume 26, Number 2, July 1992, pp. 195-202. SIGGRAPH 1992 Proceedings.

Evans *et al.* 81

K. Evans, Peter P. Tanner, and Marceli Wein, ‘‘Tablet-Based Valuators That Provide One, Two, or Three Degrees of Freedom,’’ *Computer Graphics*, Volume 15, Number 3, August 1981, pp. 91-97. SIGGRAPH 1981 Proceedings.

Fitts 54

P. M. Fitts, ‘‘The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement,’’ *Journal of Experimental Psychology*, Volume 47, 1954, pp. 381-391.

Foley *et al.* 90

James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes, *Computer Graphics: Principles and Practices*, Addison-Wesley Publishing Company, Menlo Park, California, 1990.

Herndon *et al.* 94

Kenneth P. Herndon, Andries van Dam, and Michael Gleicher, ‘‘The Challenges of 3d Interaction,’’ *SIGCHI Bulletin*, Volume 26, Number 4, October 1994, pp. 36-43.

Hudson 92

Scott E. Hudson, ‘‘Adding Shadows to a 3D Cursor,’’ *ACM Transactions on Graphics*, Volume 11, Number 2, April 1992, pp. 193-199.

Kaufman 74

Lloyd Kaufman, *Sight and Mind: An Introduction to Visual Perception; Chapter 7: The Cues to Depth*, Oxford University Press, New York, 1974.

Landy *et al.* 95

Michael S. Landy, Larry T. Maloney, Elizabeth B. Johnston, and Mark Young, ‘‘Measurement and Modeling of Depth Cue Combination: In Defense of Weak Fusion,’’ *Vision Research*, Volume 35, Number 3, February 1995, pp. 389-412.

Liang and Green 94

Jiandong Liang and Mark Green, ‘‘JDCAD: A Highly Interactive 3D Modeling System,’’ *Computers & Graphics*, Volume 18, Number 4, 1994, pp. 449-506.

MacKenzie and Buxton 93

I. Scott MacKenzie and William Buxton, ‘‘A Tool for the Rapid Evaluation of Input Devices using Fitts’ Law Models,’’ *SIGCHI Bulletin*, Volume 25, Number 3, July 1993, pp. 58-63.

MacKenzie *et al.* 91

I. Scott MacKenzie, Abigail Sellen, and William Buxton, ‘‘A Comparison of Input Devices in Elemental Pointing and Dragging Tasks,’’ in *Proceedings of ACM CHI’91 Conference on Human Factors in Computing Systems Systems*, ACM, New York, 1991, pp. 161-166.

Mackinlay *et al.* 90

Jock D. Mackinlay, Stuart K. Card, and George G. Robertson, ‘‘Rapid Controlled Movement Through a Virtual 3D Workspace,’’ *Computer Graphics*, Volume 24, Number 4, August 1990, pp. 171-176. SIGGRAPH 1990 Proceedings.

McLendon 92

Patricia McLendon, *Graphics Library Programming Guide*, Silicon Graphics Inc., Mountain View, California, 1992.

Nielson and Olsen 86

Gregory M. Nielson and Dan R. Olsen, ‘‘Direct Manipulation Techniques for 3D Objects Using 2D Locator Devices,’’ in *Proceedings of the 1986 Workshop on Interactive 3D Graphics*, ACM, New York, October 1986, pp. 175-182.

Pausch *et al.* 95

Randy Pausch, Tommy Burnette, Dan Brockway, and Michael E. Weiblen, ‘‘Navigation and Locomotion in Virtual Worlds via Flight Into Hand-Held Miniatures,’’ in *SIGGRAPH 95 Conference Proceedings*, August 1995, pp. 399-400.

Venolia 93

Dan Venolia, ‘‘Facile 3D Direct Manipulation,’’ in *Proceedings of ACM CHI’93 Conference on Human Factors in Computing Systems*, ACM, New York, April 1993, pp. 31-36.

Yoshimura *et al.* 94

Tetsuya Yoshimura, Yasuhiro Nakamura, and Masataka Sugiura, ‘‘3D Direct Manipulation Interface: Development of the Zashiki-Warashi System,’’ *Computers & Graphics*, Volume 18, Number 2, 1994, pp. 201-207.

Yuen 96

Cheuk-Mun Ada Yuen, ‘‘Self-Luminous Cursors for 3D Navigation,’’ M. Sc. Thesis, Department of Computer Science, Hong Kong University of Science and Technology, Hong Kong, 1996.

Zhai *et al.* 96

Shumin Zhai, William Buxton, and Paul Milgram, ‘‘The Partial-Occlusion Effect: Utilizing Semitransparency in 3D Human-Computer Interaction,’’ *Transactions on Computer-Human Interaction*, Volume 3, Number 3, September 1996, pp. 254-284.