Usability Analysis of 3D Rotation Techniques

*Ken Hinckley**
Microsoft Research
One Microsoft Way
Redmond, WA 98052
Tel: (425)-703-9065
kenh@microsoft.com

Joe Tullio[2], Randy Pausch[1], Dennis Proffitt[3], Neal Kassell[2]
University of Virginia
Departments of Computer Science[1], Neurosurgery[2], and Psychology[3]
{jct3u, pausch, drp}@virginia.edu

ABSTRACT

We report results from a formal user study of interactive 3D rotation using the mouse-driven Virtual Sphere and Arcball techniques, as well as multidimensional input techniques based on magnetic orientation sensors. Multidimensional input is often assumed to allow users to work quickly, but at the cost of precision, due to the instability of the hand moving in the open air. We show that, at least for the orientation matching task used in this experiment, users can take advantage of the integrated degrees of freedom provided by multidimensional input without necessarily sacrificing precision: using multidimensional input, users completed the experimental task up to 36% faster without any statistically detectable loss of accuracy.

We also report detailed observations of common usability problems when first encountering the techniques. Our observations suggest some design issues for 3D input devices. For example, the physical form-factors of the 3D input device significantly influenced user acceptance of otherwise identical input sensors. The device should afford some tactile cues, so the user can feel its orientation without looking at it. In the absence of such cues, some test users were unsure of how to use the device.

Keywords

Arcball, Virtual Sphere, 3D input devices, interactive 3D rotation, virtual manipulation, usability study, evaluation.

INTRODUCTION

With the migration of fast, cheap 3D graphics to the PC platform [26], and with 3D graphics making their presence felt on the Web through standards such as VRML [3], applications which incorporate 3D manipulation and 3D object viewing will become increasingly prevalent. In particular, orienting a virtual object to a desired view is a fundamental task, since viewing or inspecting an object is often a precursor to further manipulation.

To appear in: ACM/SIGGRAPH UIST'97 Symposium on User Interface Software & Technology

Since Chen's 1988 study (which introduced the Virtual Sphere [5]), we are not aware of any formal studies that extend Chen's results to state-of-the-art techniques. There is no formal user study to compare the Virtual Sphere with the Arcball more recently proposed by Shoemake [22][23], nor has there been a formal study with a large subject pool to explore performance advantages or disadvantages which may result from using multidimensional input devices to specify orientation. Also, unlike Chen's study, the present study includes observations of user expectations and common difficulties encountered. Using an orientation matching task based on the task employed by Chen [5], we collect performance data and investigate these issues.

One exception is a study by Ware [28], which investigates six-degree-of-freedom (6DOF) object placement and object rotation using a free-space 3D input device. For 6DOF placement, Ware found that subjects were able to make effective use of all six degrees of freedom. Ware also found that users can perform full 6DOF placements nearly as fast as the 3D orientation component of the placement by itself. Compared to our study, which takes a few measurements for a pool of 24 test users, Ware used a different research strategy, preferring a large number of measurements of only four test users. Furthermore, in the portion of his study that did look at pure 3D object rotation, Ware instructed his subjects to bias their performance towards speed, rather than accuracy.

The high cost of 3D input devices has traditionally limited their use to research or niche market applications such as head-mounted virtual reality systems, high-end animation software, or medical visualization. Free-space 3D input devices are still expensive compared to the mouse, but with the recent introduction of PC-based devices [16][2] priced

* This research was performed while the first author was affiliated with the University of Virginia, Departments of Computer Science and Neurosurgery.

near $1000, these devices are now more affordable and practical than they have ever been, and a growing number of interface designers will have the opportunity to explore the possibilities of free-space 3D input.

We do not wish to argue that any one device or technique is "best." Each interface device or input technique will excel for some tasks and languish for others [4]; the most appropriate device for an application depends on the context of tasks to be supported and the intended users. Our goal here is to collect some solid performance data for our experimental rotation matching task, so that informed design decisions can be made, and to collect some qualitative observations that will help to illustrate some strengths and weaknesses of each technique.

Interaction Techniques

Our usability analysis includes the Virtual Sphere [5], the Arcball [22], a hand-held ball-shaped orientation sensor (the “3D Ball”), and a standard magnetic orientation sensor (the “Tracker”). These interaction techniques were chosen to represent state-of-the-art techniques based on both 2D and 3D input devices, allowing us to investigate the experimental hypotheses.

The *Virtual Sphere* is a mouse-driven 2D interface which simulates a physical trackball. The virtual object is shown on the screen, and when the user clicks and drags on the virtual object, the computer interprets these drags as tugging on the simulated trackball. The virtual object rotates correspondingly. To provide the third rotational degree of freedom, a circle is drawn around the object (*figure 1*), and when the user clicks and drags in the area outside of the circle, rotation is constrained to be about the axis perpendicular to the computer screen. We will hereafter refer to this outside area of the circle simply as "the outside."

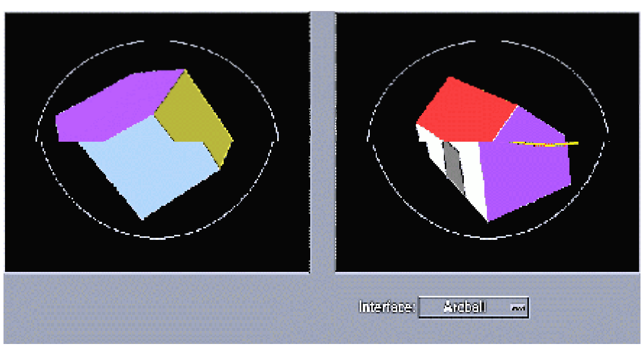

Figure 1: Screen snapshot of the experiment software. Test users rotated the house model on the right side of the screen to match the orientation shown at the left. This image shows the Arcball in use.

The *Arcball* is similar to the Virtual Sphere, but it is based upon a more mathematically elegant quaternion [21][24] implementation. It does not suffer from problems with gimbal lock or noisy data, and its implementation affords easy addition of constraint modes. It has sometimes been described as the best known 2D technique for 3D rotation. Shoemake has performed an informal comparison [22], but no quantitative data currently exist which compare the Arcball and the Virtual Sphere.

The Virtual Sphere requires the user to achieve some orientations by composing multiple rotations. In theory, with a single mouse drag the Arcball can rotate 360^{o} around any axis [23], but in practice users cannot anticipate where to start and end their dragging motion, and in effect they still must compose multiple rotations.

The *3D Ball* is a two-inch diameter plastic sphere instrumented with a magnetic tracker, which the user can rotate to manipulate the virtual object. The magnetic tracker simultaneously provides all three rotational degrees of freedom, so in principle the user never has to mentally compose rotations with this interface. However, it is not clear if users can employ coupled rotation axes effectively [5], nor is it clear if the multiple degrees of freedom result in faster, but possibly less accurate, input of orientation data.

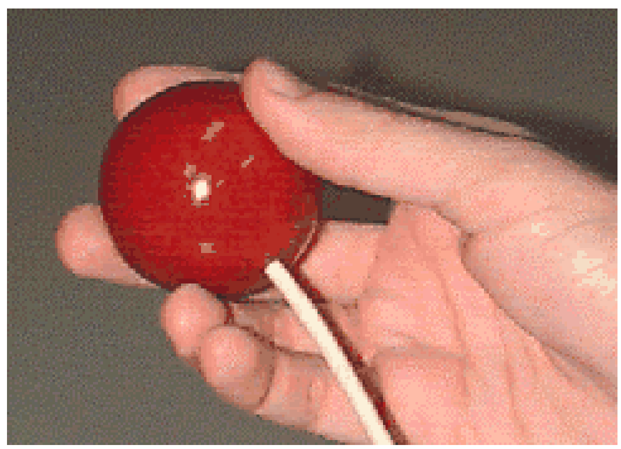

Figure 2: The 3D Ball input device.

The 3D Ball always acts as an absolute rotation controller: the orientation of the object being manipulated always matches the orientation of the 3D Ball. With the addition of a clutching mechanism (for engaging and disengaging the ball from a virtual object), it would be possible to use the 3D Ball as a relative rotation controller, by performing "ratcheting" movements [29]. The issue of absolute vs. relative rotation is an interesting secondary variable, but not a subject of this study.

Some other ball-shaped 3D input devices provide integrated buttons for clutching, ratcheting, or other functions [6][16][25], but in our experience, when buttons are integrated with the device, this can potentially interfere with manipulation (we will return to this issue in the Discussion section of this paper). Since most commercially available ball-shaped devices do include one or more buttons, we chose to construct our own 3D Ball. By excluding integrated buttons from our 3D Ball design, we can be sure that this issue will not confound our experimental results.

The *Tracker*, which also uses a magnetic orientation sensor, is identical to the 3D Ball in all regards except the physical packaging (fig. 3). This is the default form for the input device as shipped by the manufacturer [16], and as such represents the only way to use the device without designing or purchasing an alternative housing [7].

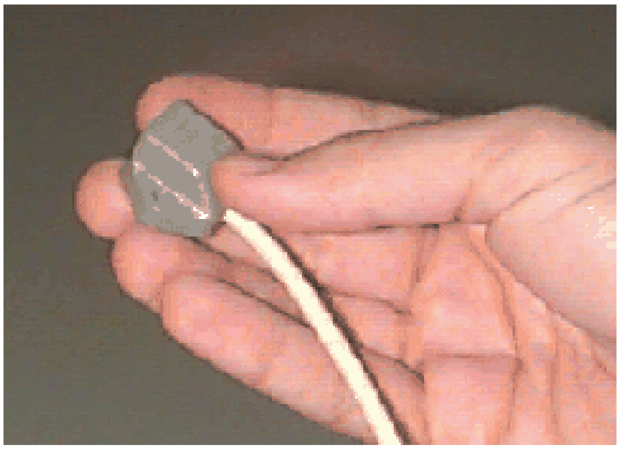

Figure 3: The Tracker 3D input device. This is the default form for the device as shipped by the manufacturer.

The Tracker has an unusual and unfamiliar shape. In our virtual reality lab [27], we have noted that with practice, experts can become quite proficient with the Tracker despite its awkward shape. It is not clear how well novice users will be able to adapt to its design.

HYPOTHESES

This study was conceived to investigate the following specific hypotheses:

H1: Users can effectively use coupled rotation axes, and integrated control of all three degrees-of-freedom for rotation will provide significantly faster input of orientation data.

Jacob [12] suggests that multiple degree-of-freedom input will be most appropriate when users think of a task's control parameters as integral attributes; we propose that 3D rotation matching is one such task. Most people are not good at mentally composing rotations, so when attempting to perform complex rotations, the separated 2D+1D control required by the Arcball and Virtual Sphere techniques should reflect this.

H2: Computer-based three-dimensional input offers the potential for faster interaction, but possibly at the expense of precision. Our hypothesis is that, at least for a 3D orientation matching task, a multidimensional input device can provide fast orientation input without necessarily sacrificing any accuracy.

H3: The physical shape (or affordances) of the 3D input device can be an important design consideration in itself.

This hypothesis arose from our previous and ongoing work with neurosurgeons to develop a free-space 3D interface for medical visualization software [9]. Surgeons could more easily learn and use this visualization interface when the 3D input sensors were presented as tools or props which helped the surgeons to reason about their tasks. The input devices took advantage of natural affordances (as discussed by Norman [15]), which can help users to know what to do just by inspecting or grasping an object or input device. The 3D Ball and Tracker used in this experiment are more general-purpose 3D input devices, yet nonetheless each communicates natural affordances which will implicitly channel user behavior; we intend to explore these issues in our analysis.

H4: The Arcball includes several apparent improvements over the Virtual sphere. As such, we expect the Arcball to outperform the Virtual sphere in terms of task performance, user acceptance, or both.

THE EXPERIMENT

Task

Our test users performed an orientation matching task based on the task employed by Chen [5]. Our goal is not to reproduce Chen's results, but rather to extend the set of interactive 3D rotation techniques that have been formally evaluated with a comparable task.

A static view of a solid-rendered 3D model of a house, at a randomly generated orientation [1], was shown on the left side of the screen (*figure 1*). Test users attempted to manipulate a second view of the house on the right-hand side of the screen to match the random orientation. Each side of the house was colored uniquely to facilitate the matching. A circle was always drawn around both images of the house to assist matching, even though the circle only was strictly necessary for the Arcball and Virtual Sphere techniques.

When test users felt that the orientations matched, they clicked a footpedal to end the trial. We used a footpedal, rather than the spacebar used by Chen [5], to keep the desk surface open for manipulation and to allow test users to use both hands (if desired) to manipulate the 3D input devices. The keyboard was removed from the desk during all experimental conditions.

After each trial, performance was rated as "Excellent!" (shortest-arc rotation less than 5.7 degrees), "Good Match!" (less than 7.6 degrees), or "Not good enough, try harder next time." The main motivation for this rating system (identical to that used in Chen's experiment [5]) was to encourage subjects to achieve a high level of accuracy; in practice, receiving a score of "Excellent" was fairly difficult. All trials (even those rated as "Not good enough") were included in our analysis. Although we instructed test subjects that speed and accuracy were equally important, it was our impression that this rating scheme tended to bias subjects to place more emphasis on the accuracy of performance.

We gave the participants as little instruction as possible with each controller. With the Arcball and Virtual Sphere, the experiment began with a practice exercise to encourage the test user to click and drag within the circle (for 2 degree-of-freedom rotation) as well as outside the circle (for the third degree of freedom). When presenting the 3D Ball or Tracker devices, we placed the device on the table in front of the test user, mentioned that "It is not a mouse, it works differently," and suggested "Try to discover how to use it." For all interfaces tested, subjects were given about 2 minutes to experiment with each device before the first trial began.

Experimental Design

A within-subjects latin square design was used to control for order of presentation effects. Test users tried all four interfaces in a single session lasting about 1 hour. Test users performed matches for 15 unique orientations with each interface, but only the last 10 of these were included in our analysis, to avoid any transient initial learning effects (our data suggest that performance was fairly stable after 5 trials; see fig. 5). There was a short break between conditions, during which we interviewed each test user about his or her impressions of the interface technique.

All subjects matched the same sequence of (randomly generated) target orientations. Thus, the total distance of rotation was counterbalanced for all subjects and devices.

Dependent variables were Time to completion and Accuracy of the match. Accuracy was measured by the shortest-arc rotation between the final user-specified rotation and the ideal matching orientation.

The data were collected on a Hewlett-Packard workstation with a 19" 1280x1024 true color display (0.28 mm^2 pixel size). The houses were rendered into 600x600 pixel windows (see fig. 1). During interactive manipulation, the display update rate was always between 22 and 25 Hertz. We used the Polhemus FASTRAK magnetic tracking system [16]; the sampling rate for the Polhemus was 30 Hz.

Test Users

Twenty-four unpaid test users (12 male, 12 female, right-handed, mean age 19.1 years) were recruited from our psychology department's subject pool. All test users had experience with the mouse, while none had any experience with 3D input devices.

RESULTS

Figure 4 shows the mean completion time and accuracy which test users achieved. The 3D Ball was 36% faster than the 2D techniques and the Tracker was 33% faster. There was little variation in the mean accuracy obtained.

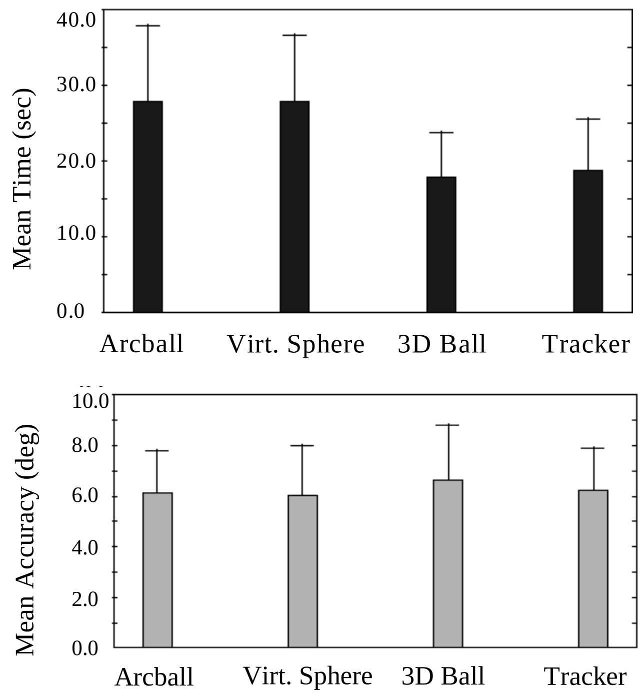

Figure 4: Mean time (top) and accuracy (bottom)

Comparing to Chen's results for the Virtual Sphere, our test users had longer times (Chen reported a mean of 17.5 seconds for complex rotations [5], while we found 27.7 seconds), but our test users were more accurate (Chen reported 8 degrees of error[1], while our test users achieved 6 degrees).

These discrepancies are probably primarily due to the differing test user populations: Chen used all males, some graduate students, and some students with experience in 3D graphics systems. For example, our male test users averaged 23.1 seconds with the Virtual Sphere (*table 3*), which is much closer to the value reported by Chen. Similarly, our sample mainly consisted of *undergraduates* who were *not* engineering students; by contrast, Chen's sample effectively selected for people who have well-developed spatial visualization skills.

It is possible that Chen's users became more proficient with the experimental task (Chen's users performed 27 trials vs. 15 trials in our study), but we have no data to support this. As shown in figure 5, performance varied little after the fourth or fifth trial. We are not sure why these two groups

[1] Chen reported accuracy in terms of the sum of squared errors between the user's rotation matrix and the rotation matrix to match. Chen indicates a squared error of 0.04 [5], which converts to 8.1 degrees.

exhibited differing accuracy levels.

A comparison with Ware's results may also be useful, although we should note that Ware studied a different task for object rotation [28], and subjects were instructed to optimize for speed at the expense of accuracy. Ware's four subjects averaged 13.4 seconds with 8.4 degrees of error, whereas in our 3D Ball condition, our test users averaged 17.8 seconds and 6.7 degrees of error.

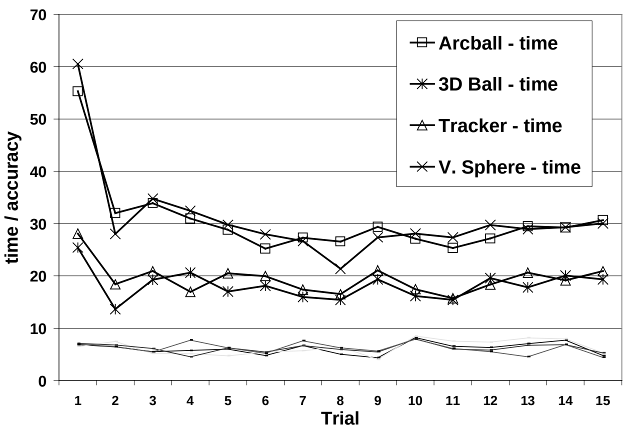

Figure 5: Effects of learning for Time and Accuracy over the 15 trials of our experiment. The two lines clustered at the top are completion time for the Arcball and Virtual Sphere; the middle lines are completion time for the 3D Ball and Tracker; the bottom cluster of lines shows learning trends for accuracy.

For the accuracy outcome measure, which results from error in matching the target orientation, our experimental task cannot distinguish between errors due to difficulty in perceiving how close the orientations match versus errors due to inadequate control of the input device. However, if errors in the control of the input device are at the level where they can hardly be perceived, presumably they have little impact on interface design. Nonetheless, this may be an interesting issue for future study.

Statistical Analysis

We performed an analysis of variance with repeated measures on the within-subjects factor of interface used, with task completion time and accuracy as dependent measures. The interface used was a highly significant factor for completion time ($F_{(3,69)} = 37.89$, $p < .0001$) but not for accuracy ($F_{(3,69)} = 0.92$, $p > .4$, n.s.).

Comparisons for completion time (*table 1*) revealed that the 3D interfaces were significantly faster than the 2D interfaces, but there was no significant difference between the 3D Ball vs. Tracker, nor was there any significant difference between the Arcball vs. Virtual Sphere.

These results strongly supports hypothesis H1, suggesting that users can perform an orientation matching task significantly faster when the rotation is presented as three integrated degrees of freedom.

Comparison	**F statistic**	**Significance**
3D Ball vs. Arcball	$F_{(1,23)} = 58.96$	p < .0001
3D Ball vs. Virt. Sphere	$F_{(1,23)} = 56.24$	p < .0001
3D Ball vs. Tracker	$F_{(1,23)} = 0.83$	p > .35, n.s.
Tracker vs. Arcball	$F_{(1,23)} = 47.31$	p < .0001
Tracker vs. Virt. Sphere	$F_{(1,23)} = 50.80$	p < .0001
Arcball vs. Virt. Sphere	$F_{(1,23)} < 0.01$	p > .95, n.s.

Table 1: Interface comparisons for completion Time

Comparisons for accuracy confirmed that there were no significant differences between any of the interfaces. This supports H2, suggesting that any accuracy differences between the interfaces are nonexistent or too small to detect with N=24 test users.

The analysis also revealed that the between-subjects factor of Sex was significant for completion time, as were the Interface X Sex and Interface X Order interactions for both completion time and accuracy (*table 2*). This indicates that we need to investigate possible biases in the data due to the Sex or Order factors.

Factors for Time	**F statistic**	**Significance**
Sex	$F_{(1,19)} = 9.69$	p < .005
Interface X Sex	$F_{(3,57)} = 3.35$	p < .03
Interface X Order	$F_{(9,57)} = 2.85$	p < .01
Factors for Accuracy	**F statistic**	**Significance**
Interface X Sex	$F_{(3,57)} = 4.79$	p < .02
Interface X Order	$F_{(9,57)} < 2.01$	p > .06

Table 2: Significant factors for the within-subjects ANOVA on factors of Sex and Order of presentation, for both time and accuracy.

Separate Analysis for Males and Females

We did not design this study to analyze Sex differences, yet our results suggest that Sex may have been a significant factor. In fact, in a pilot study where we did not carefully control for Sex, a disparity of males and females in the different Order groupings made it difficult to interpret our pilot results. Thus, as a practical matter we felt that it was necessary to control for this factor in the final study. The present study offers no data on why such a difference might exist, and our data should not be interpreted as such.

Nonetheless, our finding of slightly faster performance for male subjects is apparently consistent with the sex differences literature, which has found an advantage for males on some tasks which involve mental rotation [8]. To ensure that Sex was not a distorting factor in the final study, we performed separate analyses with the N=12 male and

N=12 female test users (*table 3*).

Males	Time (sec)	Accuracy (deg)
Arcball	22.1	6.3
Virtual Sphere	23.1	6.4
3D Ball	14.9	6.3
Tracker	15.9	6.2
Females	**Time (sec)**	**Accuracy (deg)**
Arcball	33.5	5.9
Virtual Sphere	32.4	5.7
3D Ball	20.7	7.0
Tracker	21.5	6.2

Table 3: Separate means for males and females.

For completion time, the results of the separate analysis were similar to those obtained in the combined analysis, suggesting that Sex was not a distorting factor.

For accuracy, there was a relatively small, but significant, effect for females only (*table 4*). Females were about 1 degree more accurate (*table 3*) when using the mouse-based techniques vs. the 3D Ball, but not vs. the Tracker. This suggests a minor qualification to H2, that 3D input (at least with the 3D Ball) may be slightly less accurate than 2D input for females but not for males.

Females	F statistic	Significance
3D Ball vs. Arcball	$F_{(1,11)} = 4.91$	$p < .05$
3D Ball vs. Virt. Sphere	$F_{(1,11)} = 5.02$	$p < .05$

Table 4: Significant accuracy device comparisons for females only.

Between-subjects Analysis

We also performed a between-subjects analysis using only the data from the first interface that each test user tried. Thus, the 24 test users were divided into 4 groups with 6 users each (*table 5*), resulting in an analysis which eliminates the Interface X Order effects (*table 2*) as possible factors, but also resulting in an overall less sensitive analysis.

Interface	Time (sec)	Accuracy (deg)
Arcball	30.8	5.1
Virtual Sphere	31.9	7.3
3D Ball	19.9	8.6
Tracker	21.3	7.3

Table 5: Means obtained for the first interface tried. Each mean results from only N=6 test users.

The completion time data shown above mirror the results reported in table 1, namely that the 3D interfaces are significantly faster than the 2D interfaces, but there is no reliable difference between the 3D Ball vs. Tracker, nor for the Arcball vs. Virtual Sphere.

The accuracy data suggested that test users were less accurate on the first interface tried for all interfaces, except for the Arcball (*table 6*). We can offer no explanation for this result, but the within-subjects analysis strongly suggests that any such initial accuracy differences between conditions evened out as the test user became more experienced with the task. Nonetheless, even though we only have data for N=6 test users in each group, the between-subjects analysis suggests that we may need to qualify H2: The 3D Ball may be less accurate than the Arcball when the user is first exposed to the orientation matching task.

Comparison	F statistic	Significance
3D Ball vs. Arcball	$F_{(1,23)} = 12.60$	$p < .002$
Tracker vs. Arcball	$F_{(1,23)} = 3.75$	$p < .07$
Arcball vs. Virt. Sphere	$F_{(1,23)} = 5.20$	$p < .05$

Table 6: Significant effects for between-subjects analysis of accuracy.

Subjective Ranks of Each Interface

Figure 6 shows a histogram of the subjective ranks which test users assigned to each interface. The Arcball ranks suggested a generally positive reaction, but this trend was not significant. The 3D Ball was most commonly selected as the preferred technique for the experimental task, and nobody rated it as the worst technique. Statistical analysis using the χ^2 (chi-squared) test confirmed that the 3D Ball had significantly higher ratings than the Arcball, Virtual Sphere, and Tracker techniques ($\chi^2 = 26.7$, $p < .002$).

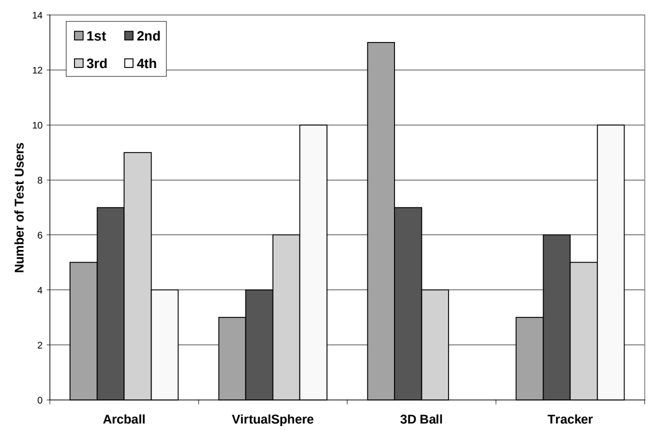

Figure 6: Subjective ranks for each interface.

This provides strong evidence in favor of H3, that the physical form factors of a 3D input device can be an important design consideration. The awkwardness of the Tracker resulted in poor subjective impressions, despite the relatively high task performance which test users were able to achieve with it. Yet the exact same input sensor packaged as a 3D Ball resulted in significantly more favorable reactions.

QUALITATIVE RESULTS

2D Techniques: Arcball & Virtual Sphere

The Arcball and Virtual Sphere shared many qualities, so we will discuss their similarities before contrasting the techniques. The techniques were generally well accepted, with many users commenting that the techniques were "pretty easy" or that they "worked really well once you learned the inside and outside of the circle."

The most common problems related to the modal distinction between the "inside" and the "outside" of the circle. During practice, we were careful to have users try both dragging inside and outside of the circle. But users see the circle as a target, and not as a border distinguishing two separate modes, so users would frequently attempt to click on the circle itself, and thus mistakenly switch between the inside and outside behaviors. Similarly, when attempting a large single-axis rotation using the outside of the circle, users would mistakenly come inside the circle, and would be surprised when the initial "outside" behavior changed to the "inside" behavior.

Several test users avoided using the outside, and would sometimes become unsure of what to do next if all the rotations matched except the rotation about the third axis perpendicular to the screen. A previous pilot study had revealed that even though this third rotation axis was available, it was a hidden feature of the interface which very few users would discover on their own. In the final study reported here, during initial practice we had each user try dragging the mouse in the outside area of the circle, yet some test users still chose to ignore this feature.

Many test users were uncertain about where to click and drag with the mouse, and once they started to drag, they were reluctant to stop. This resulted in the impression that using the mouse was "not as smooth as [the 3D techniques]- you have to start and stop a lot to click the mouse button." Some test users hesitated to make a large movement which would disturb the progress made so far. As one user commented, "When you get close, you can screw it up -- I didn't want to mess up what I already had."

Thus, test users sometimes seemed to be unsure of what effect their actions would have, and as a result they would try to plan their motions carefully. This is not a behavior we saw with the 3D techniques.

Arcball vs. Virtual Sphere

In theory, there are two primary distinctions between the behavior of Arcball and that of the Virtual Sphere: the Arcball avoids "hysteresis effects" and the Arcball uses half-length arcs. Hysteresis occurs when "closed loops of mouse motion may not produce closed loops of rotation" [22]. This means that it may not be possible to "undo" a sequence of drags by reversing their order. Half-length arcs are a property of how rotations combine, and result in a fixed C:D (Control:Display) ratio which is free of hysteresis. For example, with half-length arcs, a sweep across Arcball's inner circle moves the virtual object 360 degrees. The same sweep would only move the Virtual Sphere 180 degrees. This rotational C:D ratio is fixed by the mathematics underlying Arcball and cannot be changed without reintroducing hysteresis.

Some users did not notice these differences between the Arcball and Virtual Sphere techniques; as one test user put it, "I felt like I was doing the same thing again." Nonetheless, there was a general preference (16/24 participants) for the Arcball over the Virtual Sphere. A typical reaction was: "The Arcball is a little more responsive, and it gives you more control." Test users who preferred the Virtual Sphere often commented that they liked its slower rate of movement. The Virtual Sphere's slow movement suggested that the mouse was "pushing on a specific spot," whereas "Arcball just rotated around in different axes."

The Arcball also displays feedback arcs to illustrate how mouse movement affects rotation. These feedback arcs did not seem to be a significant advantage. The feedback arcs were often ignored or even regarded as annoying, although at least a few test users thought they helped at first. The exception was feedback arcs for the outside of the circle (for rotating about the axis perpendicular to the screen), which test users did find to be helpful.

3D Ball

Overall, test users had very positive reactions to the 3D ball technique. Typical comments were "this makes you have total control," "it was like holding the object," and "you could just turn it around instead of repeatedly clicking [with the mouse]."

Unlike the mouse-based techniques, the 3D manipulation techniques had to consider the physical form of the input device, as none of the test users had previous experience with direct 3D input. The 3D Ball's form-factors help to convey a clear message: balls are for rolling, spinning, or turning. However, this didn't always assist learning of the device. Several test users were initially convinced that the 3D Ball should be used by rolling it on the desk surface, rather than by picking it up. This "rolling" strategy was especially problematic for orientations which required the cord to point directly downward.

Although the ball shape conveys a clear message, its smooth, spherical aspect was sometimes problematic, because the surface offered few tactile landmarks or handles. This seemed to prevent some test users from forming a clear conception of the device. One test user commented that "I don't think I ever figured out how to use it-- I just wiggled it around in my fingers until I found it [the matching rotation]." The smooth surface was also somewhat slippery, making it more difficult for some test users with small hands to tumble the ball.

Most test users found the 3D Ball's cord annoying, as it would sometimes get in the way or become tangled. Some users also felt that they weren't as precise with the 3D Ball as with the mouse techniques, and some preferred using the mouse because they were used to it.

Tracker

The Tracker's physical form-factors do not convey any clear message. Test users were typically quite confused when first encountering the Tracker ("this is a very weird thing"), but they were able to adapt after several trials. One test user explained that "at first I thought it had this really dumb shape, but it turned out to be easier to hold than I thought it would be."

Several test users initially tried to use the tracker by sliding it on the desk, and users often experimented with alternative grips for holding the device in the air. Users commented that they were "unsure how to use it," that it was "an odd shape for manipulation," or that the device seemed to be too small to grasp effectively. Many test users used both hands to hold the tracker, but this did not seem to be by choice. As one test user indicated, "you almost had to use two hands" to manipulate the tracker effectively. With the 3D Ball, most test users employed one hand only.

The hindering effect of the cord was the most common verbal complaint. The weight of the cord is comparable to the weight of the tracker itself, which makes it difficult to rotate the device about its center of mass, and results in some awkward postures.

The irregular shape of the tracker caused much confusion, but it also conferred some advantages. The mounting flanges for the tracker (*figure 3*) served as handles, and gave some tactile information as to the orientation of the device. One user commented that the "ball was a big smooth object, but now I have some handles and landmarks." Three test users who struggled with the 3D Ball performed quite well with the Tracker: the tactile cues afforded by the device seemed to be essential for these individuals.[2]

DISCUSSION

Our observations from this formal study, as well as our prior informal observations from implementing several variants of 3D input devices for orientation [9][25] suggest some general design parameters which can influence how users will employ these devices. These are not well-formulated principles for design, but rather some issues intended to demonstrate how a design can implicitly channel user behavior to differing styles of interaction.

Affordances: The input device needs to suggest that it serves to orient virtual objects, while at the same time suggesting that it should be picked up, and not used on the desk surface. In our work with neurosurgical visualization [9], a miniature doll's head serves this purpose well: the doll's head just can't be manipulated effectively on the desk surface, and so strongly encourages being picked up. The 3D Ball presented in this experiment did not suggest this well enough, leading some test users to initially use it by rolling it on the desk surface.

Tactile Cues: For absolute rotation control, the device should not necessarily be completely symmetric and should have clear tactile information which indicates preferred orientation and helps to hold the device securely. In this regard, a ball-shaped device has some shortcomings. For relative rotation control, however, this issue needs to be explored further. The symmetric ball shape might be advantageous in this case, since any landmarks on the device would not correspond to the virtual object, and could be misleading. We see this as an example of how a drawback for one style of input may be a virtue for another.

Grasp: Many 3D input device designs encourage the user to hold the device against the palm at a fixed orientation. This is known as the power grasp [14] because this hand posture emphasizes strength and security of the grip. By contrast, the precision grasp involves the pads and tips of the fingers and so emphasizes dexterity and free tumbling of the input device. This issue was formally analyzed by Zhai for a 6 DOF docking task, where it was found that using the fine muscle groups emphasized in the precision grasp results in significantly faster performance [31].

The design of a 3D input device can influence which style of grasp users will choose. Integrating buttons with the input device encourages users to adopt a power grasp, holding the device in the palm, while the fingertips maintain contact with the buttons. The power grasp may require relative rotation control to avoid awkward postures, since the biomechanical constraints of the hand and arm can limit the range of rotation. Also, the muscle tension required to press or hold an integrated switch can interfere with the user's ability to manipulate the input device. Using a footpedal to separate the button from the device can be an effective alternative [9], and can help to encourage use of the precision grasp.

Device acquisition time: Our results suggested that the mouse-based techniques were slower for our experimental task. Nonetheless, in a work routine that uses the mouse frequently, the time to switch over to a 3D device, and then back to the mouse, might become a dominating factor. This could result in longer times for users to accomplish their overall goals, even though individual virtual object rotation

[2] These are the three test users who rated the Tracker as their favorite technique (fig. 6).

tasks could be performed more quickly. This illustrates why one cannot conclude from this study that "3D input is best," but that one must also consider the context of surrounding tasks and the intended users.

It may be possible to eliminate acquisition times in some cases by using a 3D device in the nonpreferred hand. This offers the possibility of working with a 3D device and a mouse at the same time [13], or with a pair of 3D devices simultaneously [9][18][19]. We have also experimented with using a touchscreen in conjunction with 3D input as a means to provide integrated 2D and 3D input in the same user interface [10][11].

CONCLUSION

Revisiting our original hypotheses, this study provides clear evidence that test users were able to take advantage of the integrated control of 3D orientation input to perform a rotation matching task more quickly than with 2D input techniques, despite years of prior experience with the mouse. More importantly, with possible slight qualifications for female test users or during initial exposure to the task, there were no statistically reliable accuracy discrepancies between any of the input techniques, demonstrating that 3D orientation input is fast without necessarily sacrificing accuracy.

The performance data did not lend any support to our hypothesis that the physical shape of the 3D input device can be an important design consideration, but the subjective ranks which test users assigned to the interfaces spoke definitively on this point. Users reported diametrically opposed impressions of input sensors which differed only in their physical housing (as shown in fig. 6).

The study did not provide evidence that the Arcball performs any better than the Virtual Sphere. Test users tended to prefer the Arcball over the Virtual Sphere, but this advantage was not significant statistically. These non-results do confirm, however, that the Arcball's possibly confusing concept of half-length arcs does not cause any obvious usability problems [22].

We believe that it would have been useful to perform standard mental rotation tests [20] on our test users after the experimental trials. We expect there would be a strong correlation between mental rotation times and task completion times. The ability to characterize test users by a standard mental rotation test would also facilitate comparison between studies.

This study has focused on manipulative techniques to achieve 3D rotation, but of course direct manipulation (whether 2D or 3D) is not the only approach to orienting and viewing objects. Even with a good interface, specification of orientation can be a difficult task for some users, so it makes sense to obviate the task when appropriate. For example, the Sketch system [30] provides a casual, inexact interface and a cleverly chosen set of heuristics which often allows the system to select views of objects based on fast, relatively imprecise 2D mouse clicks. The Jack system [17] takes a similar approach by automatically generating unobstructed views of selected objects. Additional work is needed to understand what classes of user tasks are appropriate for each technique, and to determine how to best mingle the two styles in the same interface.

We have attempted to provide a careful analysis of performance in terms of time and accuracy and to provide detailed observations of usability problems one might expect for novice test users performing tasks similar to our experimental rotation matching task. We have also attempted to synthesize specific qualitative observations from this study and our previous experience with designing 3D input devices for rotation control to suggest some tentative design parameters which can influence how users will employ these devices. Our hope is that these contributions will prove useful to the growing number of interface designers who will incorporate 3D rotation control techniques into their applications.

ACKNOWLEDGEMENTS

We would like to acknowledge the Neurosurgery department for support of our research. We thank our test users for their participation and Bob Frazier for help constructing the 3D Ball.

REFERENCES

1. Arvo, J., "Random Rotation Matrices." in Graphics Gems II, Arvo, J., ed., Academic Press, pp. 355-356.

2. Ascension Technology Corporation, PO Box 527, Burlington, VT 05402. (802) 655-7879.

3. Brutzman, D., Pesce, M., Bell, G., van Dam, A., AbiEzzi, S., "VRML: Prelude and Future," SIGGRAPH'96, Panel Discussion, pp. 489-490.

4. Buxton, W, The Pragmatics of Haptic Input. ACM CHI'90 Tutorial 26 Notes.

5. M. Chen, S. J. Mountford, A. Sellen, "A Study in Interactive 3-D Rotation Using 2-D Control Devices," Computer Graphics, 22, 4, August 1988, 121-129.

6. Chung, J.C., "A comparison of Head-tracked and Non-head-tracked Steering Modes in the Targeting of Radiotherapy Treatment Beams," Proc. 1992 Symp. on Interactive 3D Graphics, pp. 193-196.

7. Digital Image Design Incorporated, Cricket input device, http://www.didi.com/www/cricket.shtml

8. Halpern, D., Sex differences in cognitive ability. Lawrence Erlbaum Associates, 1992.

9. Hinckley, K., Pausch, R, Goble, J., Kassell, N., "Passive Real-World Interface Props for Neurosurgical Visualization," CHI'94, pp. 452-458.

10. Hinckley, K., Pausch, R., Goble, J., Kassell, N., "A Survey of Design Issues in Spatial Input," UIST'94, pp. 213-222.

11. Hinckley, K., Goble, J., Pausch, R., Kassell, N., "New Applications for the Touchscreen in 2D and 3D Medical Imaging Workstations," Proc. 1995 SPIE Conference on Medical Imaging.

12. Jacob, R., Sibert, L., "The Perceptual Structure of Multidimensional Input Device Selection," CHI'92, pp. 211-218.

13. Leblanc, A., Kalra, P., Magnenat-Thalmann, N., Thalmann, D., Sculpting with the "Ball and Mouse" Metaphor, Graphics Interface '91, pp. 152-159.

14. Mackenzie, C., Iberall, T., The Grasping Hand, Advances in Psychology 104, Stelmach, G., Vroon, P., eds., North Holland, Amsterdam, 1994.

15. Norman, D., The Design of Everyday Things, Doubleday: New York, NY, 1990.

16. Polhemus Navigation Sciences, Inc., P. O. Box 560, Colchester, VT 05446. (802) 655-3159.

17. Phillips, C., Badler, N., Granieri, J., "Automatic Viewing Control for 3D Direct Manipulation," Computer Graphics (Proc. 1992 Symposium on Interactive 3D Graphics), pp. 71-74.

18. E. Sachs, A. Roberts, D. Stoops, "3-Draw: A Tool for Designing 3D Shapes," IEEE Computer Graphics and Applications, November 1991, pp. 18-26.

19. Shaw, C., Green, M., "Two-Handed Polygonal Surface Design," ACM UIST'94 Symp. on User Interface Software & Technology, pp. 205-212.

20. Shepard, R., Metzler, J., "Mental Rotation of Three-Dimensional Objects," Science, Vol. 171, 1971.

21. Shoemake, K., "Animating Rotations with Quaternion Curves," Computer Graphics, 19 (3), 1985, 245-254.

22. Shoemake, K., "ARCBALL: A User Interface for Specifying Three-Dimensional Orientation Using a Mouse," Graphics Interface, 1992, 151-156.

23. Shoemake, K., "Arcball Rotation Control," in Graphics Gems IV, Heckbert, P., ed., Academic Press, pp. 175-192.

24. Shoemake, K., Quaternions, May 1994, available at URL: ftp://ftp.cis.upenn.edu/pub/graphics/shoemake/quatut.ps.Z.

25. Stoakley, R., Conway, M., Pausch, R., "Virtual Reality on a WIM: Interactive Worlds in Miniature," Proc. ACM CHI'95, pp. 265-272.

26. Torborg, J., Kajiya, J., "Talisman: Commodity Realtime 3D Graphics for the PC," Computer Graphics (Proc. SIGGRAPH'96), pp. 353-363.

27. University of Virginia user interface group. http://www.cs.virginia.edu/~uigroup/

28. Ware, C., "Using Hand Position for Virtual Object Placement," Visual Comp., 6 (5), 1990, 245-253.

29. Colin Ware, D. R. Jessome, "Using the Bat: A Six-Dimensional Mouse for Object Placement," IEEE Computer Graphics and Applications, November 1988, 65-70.

30. Zeleznik, R., Herndon, K., Hughes, J., "SKETCH: An Interface for Sketching 3D Scenes," Proc. SIGGRAPH'96, pp. 163-170.

31. Zhai, S., Milgram, P., Buxton, W., "The Effects of Using Fine Muscle Groups in Multiple Degree-of-Freedom Input," Proc. ACM CHI'96, pp. 308-315.