Comparison of relative (mouse-like) and absolute (tablet-like) interaction with a large stereoscopic work-space

Melinos Averkiou and Neil A. Dodgson

The Computer Laboratory, University of Cambridge
15 J J Thomson Avenue, Cambridge, UK CB3 0FD

ABSTRACT

We compare two different modes of interaction with a large stereoscopic display, where the physical pointing device is in a volume distinct from the display volume. In *absolute* mode, the physical pointer's position exactly maps to the virtual pointer's position in the display volume, analogous to a 2D graphics table and 2D screen. In *relative* mode, the connection between the physical pointer's motion and the motion of the virtual pointer in the display volume is analogous to that obtained with a 2D mouse and 2D screen. Both statistical analysis and participants' feedback indicated a strong preference for absolute mode over relative mode. This is in contrast to 2D displays where relative mode (mouse) is far more prevalent than absolute mode (tablet). We also compared head-tracking against no head-tracking. There was no statistically-significant advantage to using head-tracking, however almost all participants strongly favoured head-tracking.

Keywords: stereoscopic, mouse, tablet, interaction, head-tracking

1. INTRODUCTION

We consider a situation where a designer is working with a large stereoscopic display. In our case, we used a display with 100" diagonal. We conducted experiments to compare ways of interacting with such displays. With displays of this size, it is impossible for the designer to use a pointing device physically embedded in the perceived volume. We therefore investigated how best to interact when the pointing device is in a volume distinct from the drawing space. This is exactly analogous to the situation we face on 2D displays where we use a mouse or a tablet that operates in a separate 2D space, the table-top, which is remote from the 2D display.

We investigated two modes of interaction, which are 3D analogues to the 2D tablet and the 2D mouse. In *absolute mode*, the position of the pointing device in space exactly maps to the pointer's position in the perceived volume, as it would do in 2D for a tablet. In *relative mode*, the movement of the pointing device corresponds to movement of the pointer in the perceived volume, but the absolute positions do not correspond. This is analogous to a 2D mouse where the linkage between pointer position on screen and mouse position on the tabletop can be "frozen" by lifting the 2D mouse off the tabletop. We implemented a similar "freezing" mechanism for the 3D pointer. A secondary aim was to investigate whether head-tracking affected a person's performance in 3D modelling tasks, given that the display was already stereoscopic. There is evidence elsewhere that stereoscopic rendering significantly influences performance in 3D pointing tasks[1–3] and that head-tracking has a significant effect on 3D manipulation tasks.[4]

Our study suggests that absolute positioning should be preferred to relative positioning for 3D interfaces, both in general and for 3D modelling in particular. This is in contrast to 2D interfaces where the mouse (relative positioning) is favoured over the tablet (absolute positioning) for many tasks, including much 3D modelling work. The study also indicates that head-tracking is not an essential feature for the 3D tasks attempted. This contradicts previous work[4] that suggests a significant effect of head-tracking for similar tasks.

Further author information: send correspondence to N.A.D., e-mail: nad@cl.cam.ac.uk, Telephone: +44 1223 334417

Averkiou, M & Dodgson, NA, "Comparison of relative (mouse-like) and absolute (table-like) interaction with a large stereoscopic work-space", Stereoscopic Displays & Applications XXII, Woods, Holliman, Dodgson (eds.), *Proc. SPIE* **7863,** paper 7863-41, 2011

DOI: http://dx.doi.org/10.1117/12.??????

2. BACKGROUND

Our project investigated 3D interaction techniques based on absolute and relative pointer positioning, analogous to existing 2D interaction techniques.

2.1 Relative positioning

The most obvious example of relative positioning for 2D interfaces is the mouse. The mouse detects relative motion using a mechanical ball or light reflections. It transmits the direction and the amount of motion, which are then used to move a 2D pointer on the graphical user interface. It is a relative positioning device, as it does not report its absolute position according to a reference frame, but instead reports the direction and amount of its motion on a surface. Therefore, there is no direct correspondence between the mouse position and points on the computer screen.

In a limited workspace, moving the mouse cannot move the 2D pointer across the entire screen. In this case one can pick up the mouse and take advantage of the third physical dimension to move it to a new location from which to continue the on-screen motion. While the mouse is airborne, the 2D pointer does not move at all. Extending this modality to three dimensions is not trivial because there is no fourth spatial dimension to use to “freeze” the 3D pointer at a specific location. We chose to use a button to freeze and unfreeze the 3D mouse. This is not as obvious as lifting a 2D mouse off the table but is necessary in the absence of a fourth dimension.

2.2 Absolute positioning

An example of absolute positioning for 2D interfaces is the digitizing tablet. It is controlled using a pen or stylus, usually equipped with buttons. The tablet contains hardware that detects the absolute location of the tip of the stylus on the tablet’s pad and transmits this to the graphical user interface. It is an absolute positioning device because each location on the tablet’s pad corresponds to a location on the computer screen.

MacKenzie, Sellen, and Buxton[5] found that the tablet offered at least as good accuracy as the mouse for pointing and dragging tasks. They felt that the tablet had the potential to perform as well as the mouse in direct manipulation systems, and to out-perform the mouse when user activities include drawing or gesture recognition.

2.3 Physically-coincident *vs* physically-distant manipulation

Both the mouse and the digitizing tablet are used in direct manipulation modelling interfaces, but are physically distant in their operation since the virtual modelling space is rendered on the computer screen while the mouse or the tablet pen are elsewhere: on the table top or on the tablet pad. Therefore, the input space and the output space are distinct. The alternative is physically-coincident manipulation, exemplified by the touch screen, which combines 2D absolute positioning input and 2D output in one device, making the manipulation of objects even more direct than interfaces based on the mouse or the tablet.

There are 3D analogues of the touch screen. For example, there are systems that use a *SensAble Phantom* that is either behind a half-silvered mirror[6] or suspended above a display screen[7,8] to co-locate the manipulation device and the 3D display volume. This has two limitations: the manipulated 3D volume must lie entirely in front of the display screen, and the manipulated 3D volume is limited in depth. To view the 3D content, observers must be able to disengage accommodation and convergence. This can be comfortably done up to a limit of around 25% of the distance from the screen to the observer. Because the manipulation volume must lie within arm’s reach, the usable volume is limited to a depth of perhaps 20 cm, being about one quarter of the maximum stretch of the observer’s arm plus the Phantom’s stylus.

By contrast to these systems, we wanted to experiment with a large 3D workspace, where the space is entirely behind the screen. Observers are able to handle a greater range of depth, behind the screen, than is possible in front of the screen. This gives a much larger workspace than using the space in front of the screen, and avoids the problem of the observer having to be within arm’s reach of the screen. However, the manipulation device is now physically-distant from the space on which it operates. We wished to investigate whether relative (mouse-like) or absolute (tablet-like) interaction would be preferable in such a situation. In addition, we decided to investigate whether head-tracking added any significant benefit in stereoscopic viewing for 3D tasks. We decided not to investigate monoscopic viewing, as this would double the number of conditions that needed to be tested and there are already a wide range of experiments that have investigated the benefits of stereoscopic against monoscopic viewing.[1–3]

(a) (b) (c)

Figure 1. (a) A Vicon infrared camera with its circular infrared LED array. (b) The custom made pointing object (length 70 cm) with multiple retro-reflecting spheres. (c) The head-tracking hat, with multiple retro-reflecting spheres, and a set of polarised glasses.

3. METHOD

We describe the equipment and the experimental design.

3.1 Equipment

A wide range of equipment was used. The main components were a Vicon optical tracker to track the observer's head and the pointing device; a Wii Remote (Wiimote) to provide "mouse buttons"; a stereoscopic projector and a silver screen.

Optical motion tracking was provided by a Vicon MX motion tracking system (Figure 1(a)). It uses infrared cameras to detect the position of small infrared retro-reflecting spheres in 3D space. Proprietary software allows one to select multiple retro-reflecting spheres and instruct the system to track them as a rigid body. The centre and the orientation of the rigid body is then made available to user software. It is straightforward to track multiple rigid objects by mounting several infrared reflectors on each object, and instructing the system to treat each as a separate rigid object. Two objects were tracked: a hat worn on the observer's head (Figure 1(c)) and a 3D pointer.

We needed to implement a six DoF input device that would be used as a 3D pointer, analogous to how both the mouse and the stylus of a digitizing tablet act as a 2D pointer in 2D interfaces. We considered using a Wiimote for this, but the Wiimote had insufficient accuracy to be usable as the absolute positioning device, and thus could not be used. We constructed a pointing object from doweling, cardboard, and duct tape (Figure 1(b)). Vicon infrared reflectors were attached to the object to provide robust tracking of its position and orientation. It was deliberately designed to resemble a wand or pointer so people would be accustomed to the idea of using it for pointing at and picking things.

A 3D physical pointer on its own, without any control buttons, is insufficient for modelling. We used a Wiimote to provide all necessary controls. The operator used the 3D pointer in their dominant hand, with the Wiimote in their other hand. One button was used to select and deselect objects, another button was used to freeze and unfreeze the pointer in relative (mouse) mode.

Stereoscopic display was provided by a DepthQ 120 frame per second (FPS) projector at 1280×720 resolution. An alternating circularly-polarised filter was mounted in front of the projector to provide the correct polarisation for alternate frames. The projector was driven by a nVIDIA Quadro FX graphics card. The screen was 230×160cm and the stereoscopic space was about 200cm deep, entirely behind the screen. The manipulation space was centred about 3m from the screen, and was the same physical size as the perceived 3D volume behind the screen.

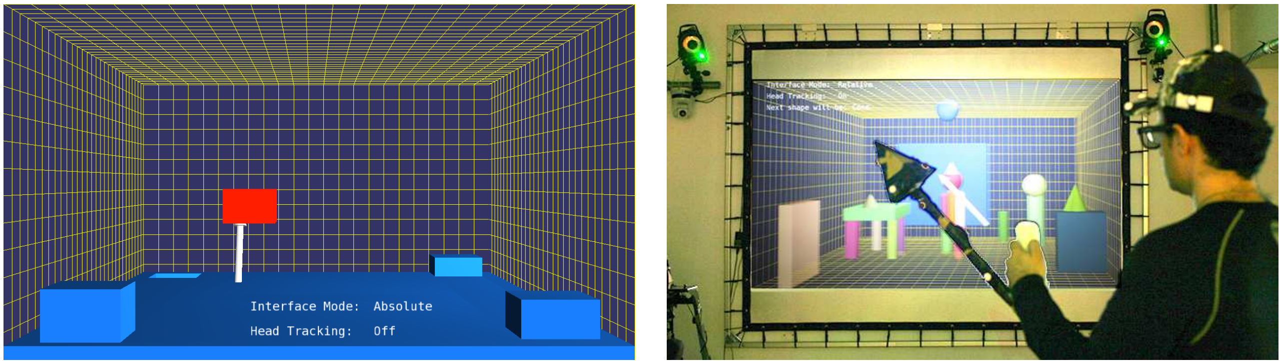

(a) (b)

Figure 2. The two tasks. (a) Screenshot of the peg-board task. (b) The first author interacting with the 3D modelling task. He holds the pointer in his right hand, and the Wiimote in his left. The hat and pointer are tracked by the Vicon system. Because the experiment was conducted in low lighting, this photograph has been enhanced: the luminance of everything other than the screen has been increased.

3.2 Experiment design

We compared the performance of absolute and relative interaction techniques and the effect of head-tracking on 3D modelling tasks. A four-way, within-subjects experiment was designed for this purpose. The four experimental configurations that were used are:

1. Absolute positioning interaction with head-tracking.
2. Relative positioning interaction with head-tracking.
3. Absolute positioning interaction without head-tracking.
4. Relative positioning interaction without head-tracking.

To perform an accurate and comprehensive comparison between the effects of the different experimental configurations, sufficient quantitative, as well as qualitative, data had to be collected. Two tasks were designed: one to collect quantitative data, through timing and precision measurements, and one to capture qualitative data, through a questionnaire. A pilot study, with two participants, was used to refine the initial experimental design.

3.2.1 Peg board task

The first task captured quantitative data about the performance of the participants in simple peg-in-hole assignments. This type of task has been used previously[3] and is similar to other docking-type tasks[1,2,9–11] used for measuring 3D pointing accuracy and speed. In our case, the peg-in-hole task was designed to combine both object selection and object manipulation, two of the universal interaction tasks in 3D interfaces.[12]

Speed and accuracy were chosen as the metrics that best characterize the performance of a participant. To gather sufficient data, the peg board was designed with five holes and five pegs that fit into the holes (Figure 2(a)). The five pairs of pegs and holes provide five speed and five accuracy measurements for every configuration of the experiment, giving a total of twenty speed and twenty accuracy measurements for every participant. This number was a trade off between gathering sufficient data and ensuring that the participants did not spend too long on the experiment.

The peg board was made horizontal as this forced participants to move pegs in all three dimensions, especially in depth. All the holes had the same size. The pegs were also of the same size and they fitted exactly in the holes, though, unlike physical pegs, it was not essential for the participant to get the peg exactly in the hole. The location for every hole and the initial location of its corresponding peg were carefully selected so that the distance covered and thus the time needed in order to move one peg to a hole would be the same in average across the five peg-hole pairs.

The task begins immediately after pressing the *A* button on the Wiimote. The first peg-hole pair appears. The participant is required to pick up the peg by pointing the 3D pointer to it and pressing the *A* button on the Wiimote. The participant then moves the peg and puts it into the hole, pressing *A* again to release it. The next peg-hole pair appears only if the distance of the peg to the hole is below 50 cm, in order to accommodate accidental presses of the button.

Accuracy is recorded as the Euclidean distance in the xz (horizontal) plane from the centre of the hole to the centre of the peg. The 3D Euclidean distance is not needed as we are not interested in how far into the hole the peg is pushed: only in how accurately it fits in the hole.

3.2.2 3D modelling task

The second task captured qualitative data. Participants were asked to create a simple object using a 3D modelling interface (Figure 2(b)). This task is closer to the conventional process of 3D modelling and was chosen to allow us to get feedback from the participants about their experience of actual 3D modelling. Liang and Green comment that existing 3D modelling systems have low efficiency of interaction owing to the need to use 2D input devices.[13] We wanted to get qualitative feedback on the differences between our 3D input modalities.

The task required the participant to draw a simple table using the 3D modelling application. For the purposes of this experiment a table was deemed to be four legs, each touching the floor, and a table top, sitting on top of the four legs. The challenges were to get the legs all of the same length, to get all of the legs to touch the floor, and to get the table top to sit on top of all four legs, without any leg poking out the top, any gap between leg and table top, or any leg not supporting the table top. Anything reasonably close to what was desired was deemed acceptable. Participants were allowed to create objects and move them, but not to rotate them. The software was simple and restricted in operation. Participants could only draw axis-aligned parallepipeds. They would press button *B* to start drawing, and *B* again to stop drawing. They could pick up an object by pressing button *A*, and press *A* again to let go of the object. There was no ability to delete, undo, or group objects. There was no restriction on the time allocated to the participant for creating the table, however it was reasonable to expect the task to last less than five minutes.

A short questionnaire was used. It was divided into two parts. The first part used questions with five point Likert-scales, based on Nielsen's attributes of usability.[14] For each of the interaction techniques, three questions solicited the participants' opinion of its efficiency, learnability, and ease-of-use. The second part used five questions with a forced two-way choice. Two questions solicited an opinion as to whether head-tracking was beneficial and whether it was preferred. The other three questions assessed which of the two interaction techniques was preferred, which was easiest to use, and which was more comfortable for extended use.

3.3 Participants and procedure

Ten people were recruited for the user study, five were computer science students and five were from different backgrounds. Three were female and seven were male. All participants were below thirty years old and all were right-handed. Three participants had prior experience with 3D modelling. The participants' ability to view stereo was tested using a 3D random dot pattern designed so that it would be obvious whether or not the participant had stereoscopic vision. All participants could see stereo well.

After explaining the purpose of the experiment, every participant was asked to sign an informed consent form. Then, the first task was introduced and the controls were explained to the participant. We wanted minimal training for the first task so the participant was only shown the first peg-hole pair.

The participant was then asked to perform the peg board task as fast and as accurately as possible. A random ordering of the four experimental configurations was used to compensate for learning effects. After running the task in all four configurations, the participant was shown the controls for the modelling task and was given five minutes to use the application to draw some shapes. It was important for the participant to test all four configurations and form an opinion for each one. The participant was then asked to build the table in the 3D modelling task in all four configurations.

Finally, the participant was asked to complete the questionnaire. During the course of the experiment, the behaviour, the reactions and any comments of the users were recorded on paper notes. A short open-ended interview about the overall experience of the participant was held after the participant had completed the questionnaire. The measurements from the peg board task, the questionnaire answers and the draft notes were then analyzed.

Table 1. ANOVA (Two-factor with replication) of users' speed. Please note that, as reported in the main text, four of the 200 time measurements were invalid. To compensate for this, we ran two alternative ANOVAs. In the one reported below, those invalid measurements were replaced with zeros (i.e., biasing the data away from significance). In the other analysis, those invalid measurements were replaced by the average time from the other four runs for that user on that configuration (i.e., biasing the data towards significance). In that other analysis, the F values were increased to 18.5 (Users), 64.8 (Configurations), and 3.63 (Users $\times$ Configurations).

Source of Variation	*SS*	*df*	*MS*	F	p *value*	F *critical*
Users	3641.208	9	404.579	6.927	2.1×10^{-8}	1.939
Configurations	6226.431	3	2075.477	35.534	1.18×10^{-17}	2.661
Users $\times$ Configurations	2798.398	27	103.644	1.774	0.016	1.556
Within	9345.241	160	58.408			
Total	22011.28	199				

Table 2. Two tailed t-tests of users' speed between pairs of experiment configurations. Key: HT (head-tracked), non-HT (not head-tracked), abs (absolute, tablet-like, positioning), rel (relative, mouse-like, positioning).

Constant factor	*Configuration Pair*		p *value*
head-tracked	1–2	abs *vs* rel	1.42×10^{-8}
absolute	1–3	HT *vs* non-HT	0.462
	1–4	abs HT *vs* rel non-HT	1.22×10^{-11}
	2–3	rel HT *vs* abs non-HT	1.16×10^{-6}
relative	2–4	HT *vs* non-HT	0.354
non-HT	3–4	abs *vs* rel	3.13×10^{-9}

4. RESULTS

Statistical analysis was performed on the speed and accuracy measurements. Speed was measured by the time needed to put a peg into a hole: a smaller value thus means greater speed. Accuracy was defined as the Euclidean distance in the xz (horizontal) plane from the centre of the hole to the centre of the peg, a smaller value thus means greater accuracy.

Twenty speed and twenty accuracy measurements were collected for each user. A total of 200 speed and 200 accuracy measurements were collected. Four speed measurements had to be rejected from the statistical analysis because they were affected by unforeseen technical problems that made them invalid, leaving 196 speed measurements and 200 accuracy measurements.

4.1 Speed

A two-way analysis of variance (ANOVA) was conducted for all the speed measurements (Table 1). There is a significant difference between the mean speed for at least one pair of the ten users over all four configurations of the experiment ($F(9, 160) = 6.93, p < 0.05$). There is also a highly significant difference between the mean speed for at least one pair of the four experiment configurations over all ten users ($F(3, 160) = 35.5, p < 0.05$). This means that at least two of the experiment's configurations had a significant effect on the mean speed of the users. Lastly, the interaction of users and experiment configurations had a significant effect on the mean speed ($F(27, 160) = 1.77, p < 0.05$). This means that both the choice of user and the choice of experiment configuration can affect the mean speed at the same time.

These results mean that there is a significant effect of the experiment's configuration on the mean speed of the users. Two-tailed t-tests between all six possible pairs of the experiment's configurations were conducted to ascertain which particular pairs had significant differences (Table 2). There is a significant difference ($p < 0.05$) between all pairs that compare absolute positioning with relative positioning, but there is not a significant difference for the two pairs that do not change the positioning method. Considering these significant differences and the mean speed for each configuration of the experiment (Table 3), we draw the following conclusions:

1. Absolute positioning is significantly faster than relative positioning. This comes from the fact that the mean speed (time to completion) between absolute and relative positioning interaction was significantly different, as illustrated

Table 3. Mean speed and accuracy of users for each configuration of the experiment

Configuration	*Mean speed (sec)*	*Mean accuracy (cm)*
1. Absolute—Head-tracking	14.065	6.197
2. Relative—Head-tracking	24.823	6.361
3. Absolute—No head-tracking	14.645	6.783
4. Relative—No head-tracking	25.985	6.831

Table 4. ANOVA (Two-factor with replication) of users' accuracy

Source of Variation	*SS*	*df*	*MS*	*F*	*p value*	*F critical*
Users	0.105	9	0.012	4.152	7.92×10^{-5}	1.939
Configurations	0.004	3	0.001	0.422	0.738	2.661
Users $\times$ Configurations	0.045	27	0.002	0.599	0.940	1.556
Within	0.449	160	0.003			
Total	0.603	199				

by t-test p-values for configuration pairs 1–2 and 3–4. In addition, the mean speed for absolute positioning was significantly less (thus faster) than the mean speed for relative positioning in both of these configuration pairs.

2. Head-tracking does not have an effect on speed since the mean speed between head-tracking and non-head-tracking modes was not significantly different, as illustrated by t-test p-values for configuration pairs 1–3 and 2–4.

4.2 Accuracy

A two-way ANOVA was conducted for all accuracy measurements collected from the peg board task (Table 4). There is a significant difference between the mean accuracy for at least one pair of the ten users over all four configurations of the experiment ($F(9, 160) = 4.15, p < 0.05$): so users do differ significantly in accuracy. In contrast to the speed measurements however, there is no significant difference between the mean accuracy of any pair of experiment configurations ($F(3, 160) = 0.422$, n.s.). This means that none of the experiment configurations had an effect on the mean accuracy of the ten users. Finally, the interaction of users and experiment configurations also did not have any significant effect on users' mean accuracy. This leads to the conclusion that neither the choice between absolute positioning and relative positioning nor the choice between head-tracking and no head-tracking has any significant effect on the users' mean accuracy.

Figure 3 graphs the speed and accuracy results. Figure 3(left) shows the difference in median time for absolute positioning (around 13 seconds) and relative positioning (around 23 seconds), indicating the magnitude of the significant different identified by the ANOVA. Head-tracking was shown to have no significant effect on speed, and this can be informally observed in the graphs.

Figure 3(right) illustrates the lack of differences in accuracy between the different conditions. All configurations had a mean accuracy between 6 and 7 centimetres. As shown by the ANOVA, none of the values for mean accuracy was significantly different than the others.

Finally, Figure 4 shows the lack of correlation between speed and accuracy across all the users for each configuration of the experiment. There is no correlation between speed and accuracy for any configuration, as the trend lines for each configuration are roughly horizontal. There is a small trend towards better accuracy as speed decreases in the relative positioning mode but this is not significant. Also, relative positioning has a more dispersed sample of speed measurements than absolute positioning whose sample is highly clustered between 10 and 20 seconds.

4.3 Questionnaire results

The first part of the questionnaire asked participants about the ease of learning, ease of use, and perceived efficiency of the two interaction modes (Figure 5). The majority of participants rated the absolute positioning technique as easier to learn than the relative positioning technique. The same is true for efficiency, as 80% of the participants ranked the efficiency of the absolute positioning technique as 4 or 5, while most participants ranked the efficiency of the relative positioning

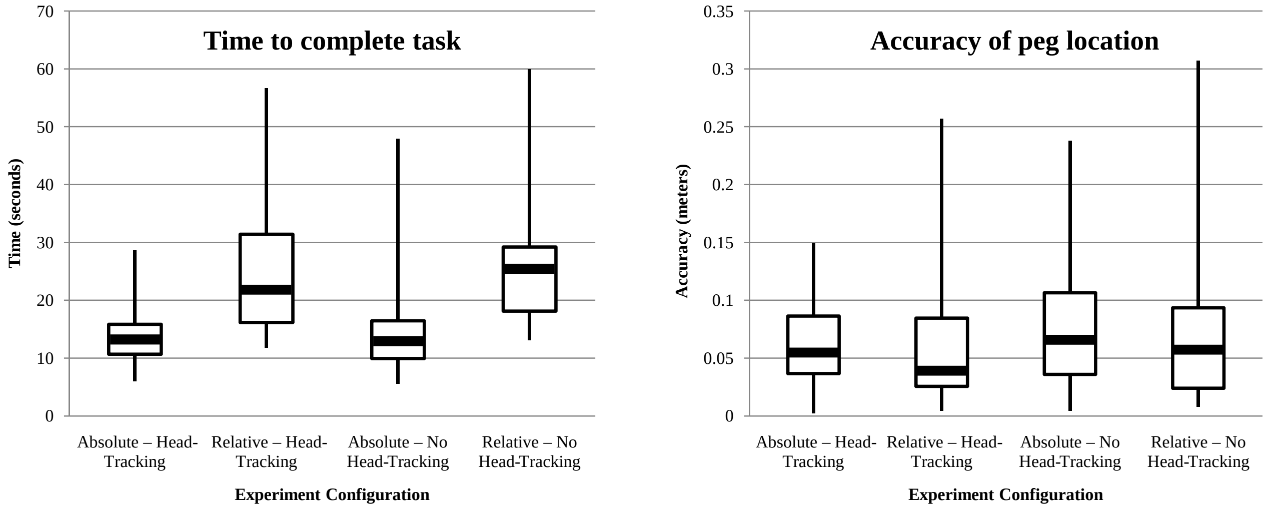

Figure 3. Graphs of (left) time to complete task and (right) accuracy. The box shows upper and lower quartiles. The heavy bar shows the median. The top and bottom of the whiskers show maxima and minima.

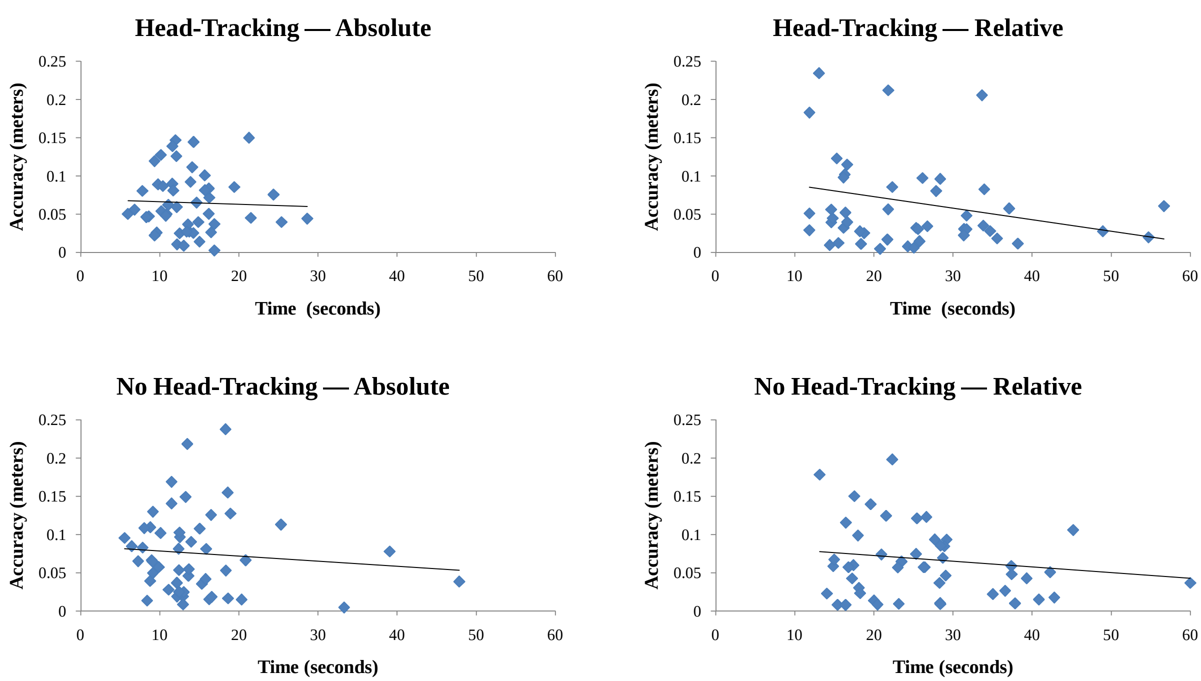

Figure 4. Accuracy as a function of time for all configurations.

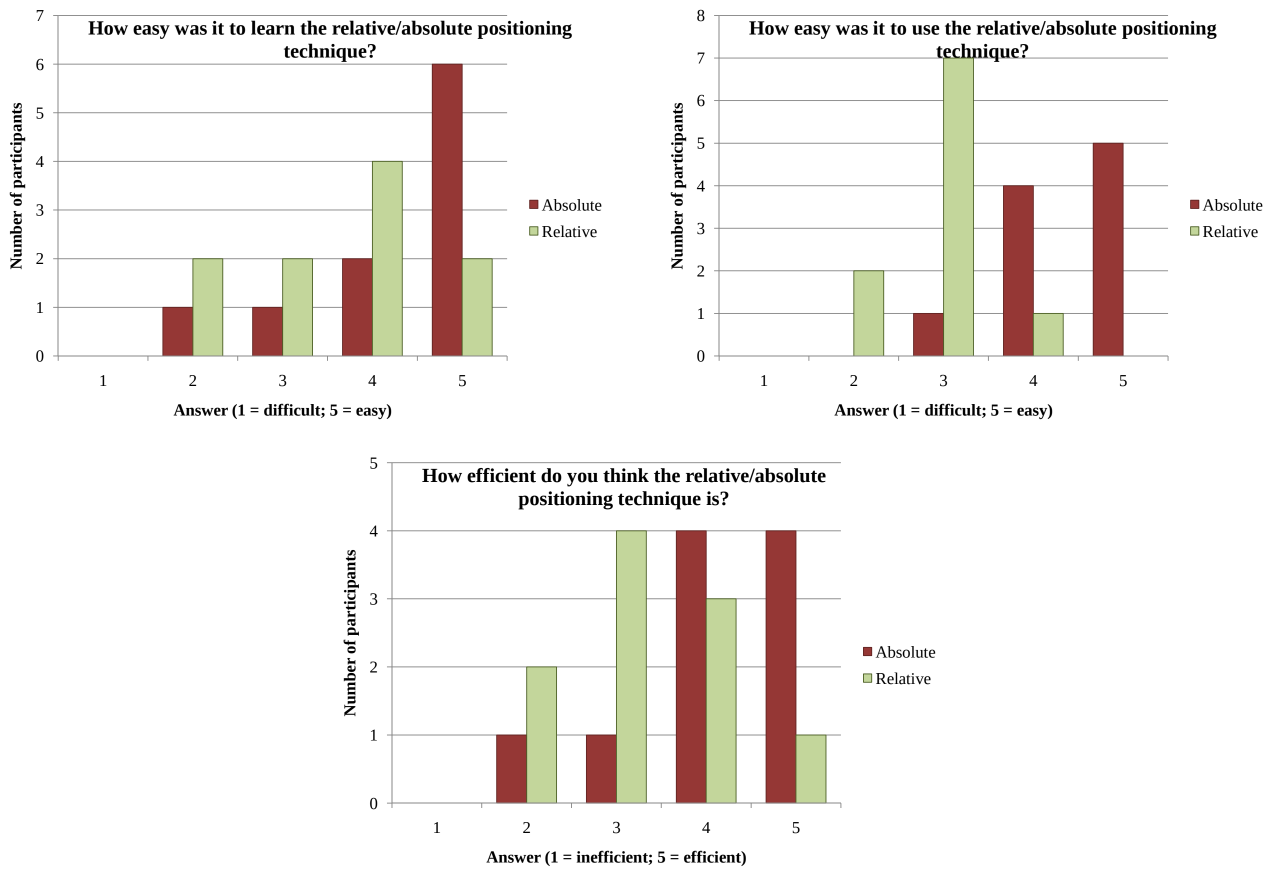

Figure 5. Answers to the questions about ease of learning, ease of use, and efficiency of the two interaction modes.

technique as 3 or 4. Most participants thought that absolute positioning was easier to use than relative positioning. Nine participants ranked the absolute positioning technique as 4 or 5, while almost all participants ranked the relative positioning technique with a 3.

The second part of the questionnaire consisted of five two-way selection answers. Two concerned head-tracking and three concerned the interaction technique. Nine of the ten participants preferred head-tracking and all of the ten thought it had helped them achieve better results, although the statistical analysis does not support this. When asked which mode they would prefer, nine out of ten found absolute mode easier to use, and eight out of ten said that they would prefer using the absolute mode, even for extended use, despite it being more active and therefore potentially fatiguing.

5. DISCUSSION

The results can be summarised in five conclusions:

1. The absolute interaction technique is significantly faster than the relative interaction technique in the peg board task.
2. The absolute and the relative interaction techniques had no significant difference in accuracy.
3. Head-tracking does not significantly affect accuracy or speed.
4. Participants preferred the absolute interaction technique.
5. Participants preferred head-tracking.

The first and second conclusions indicate that the absolute interaction technique provides better performance than the relative interaction technique for pointing and manipulation in this 3D interface. This is because it was shown to be faster than but just as accurate as the relative interaction technique.

The third conclusion suggests that head-tracking is not an essential feature for modelling in a 3D interface since it provides no significant improvement in performance. This means that the cost and effort of using it for a 3D interface can be avoided. More research is needed before this implication can be proved beyond doubt, since previous work[4] suggests that head-tracking does help.

The last two conclusions were corroborated by the observations of participants during the experiment and the insights that were gained from the experience of conducting the experiment. The first two conclusions combined with the fourth seem to suggest that absolute positioning should be preferred over relative positioning in 3D interfaces for modelling since it is better performing and people prefer it.

The last conclusion seems to contradict the third. It is surprising that head-tracking did not affect speed and accuracy since most users thought it had helped them and they had taken advantage of its benefits. The implication of this is that it becomes a matter of design choice whether one wants to satisfy user preference at the expense of tracking the user's head and thereby incurring higher processing and hardware costs.

It is worth noting how well people performed with no training, in contrast to at least one previous study[1,2] in which users took part in lengthy experiments comprised of training and tests before any actual data could be collected. An investigation into the effect of training would be interesting, as many people find 2D mouses difficult to use at first, and it may be that people would become more comfortable with the 3D mouse-like method after a reasonable amount of practice.

5.1 Internal validity of the experiment

The user study was designed so that it was protected against common factors that may jeopardize the internal validity of its results. Of the eight common extraneous variables affecting internal validity,[15] only testing was considered a threat. This is the effect that repeated testing may have on a subject's performance and it is a common pitfall in within-subject designs. This factor was alleviated by randomizing the execution sequence for the four experimental configurations so that training effects would not be observed. Maturity was another factor that was considered, but the experiment was short and the number of experiment runs was small, therefore unlikely to be a problem.

5.2 External validity of the experiment

We can reasonably expect that the results of the experiment and their implications can be generalized to other users. This is because the participants were chosen to reflect both computer science and non-computer science backgrounds and they included both males and females. Regarding the performance metrics, it can be argued that speed and accuracy are objective indicators for the performance of an interaction technique or device as they have been used extensively in the literature.[1,2,9–11,16] The peg board task was fair to all experimental configurations and was carefully selected so the results collected could be generalized for pointing and manipulation in 3D interfaces. It is a known evaluation task that has been used in the past[3] and it is an excellent way of measuring performance, which was the goal of the study. In this case it was adjusted to include both the pointing and the manipulation elements of 3D interaction.

5.3 Threats to the validity

One could argue that the choice of using a button for freezing the pointer is what caused the poor performance in the relative positioning technique. However, the choice of using a button was not random but was taken after careful consideration of the alternatives. Using a gesture of the wand as an alternative could seem faster and easier for the user, but it has two problems: A gesture moves the pointer, which will move the object to which it is attached, therefore preventing any sort of fine positioning. Furthermore, users may have difficulty in reliably making the gesture. A different solution would be to track the users' hands rather than a pointer. Unfortunately, these are more difficult to track precisely since they can change shape and they will be less accurate than holding a real-life pointer with a tip. However, it should be noted that at least a couple of the users mentioned that the relative pointing was hard to get used to and seemed unnatural.

People were observed having difficulty in coordinating both hands for pressing buttons and moving the pointer. It could be well argued that the wand and the Wiimote should be integrated in one physical device, and in fact at least one of the users mentioned that. We would have preferred this ourselves, but were limited in the equipment available. We believe that most, if not all, users performed well when using this bi-manual interaction. Another user-study would be needed to ascertain whether bi-manual or single-handed manipulation is preferable. Note that the principal difference in manipulation between the two interaction techniques is that relative positioning required two buttons rather than one.

It is possible that the relative positioning method would be preferred for fine work or for work in a stereoscopic space larger than the volume in which the pointer can be moved. We did not test either of these conditions in this experiment.

Possibly the biggest threat to the validity of the head-tracking results was the fact that participants were asked to stay in one location while using relative positioning, and were therefore unable to take full advantage of the motion parallax depth cue provided by head-tracking. This could invalidate the conclusion that head-tracking has no effect on speed or accuracy, at least when combined with relative positioning. However, there is still the result that people clearly moved sufficiently to prefer the head-tracked conditions over the non-head-tracked conditions.

6. CONCLUSIONS

Statistical analysis showed that absolute positioning outperformed relative positioning in speed ($p < 0.05$), but that there was no significant difference in accuracy. Head-tracking had no effect on either speed or accuracy. Qualitative data supported the statistical analysis regarding the interaction techniques, as participants strongly favoured absolute positioning to relative positioning. However, the qualitative data collected by the questionnaires and discussions with participants indicates that almost everyone preferred head-tracking and believed it enhanced their performance, even though there is no statistical evidence to support this.

The study suggests that absolute positioning should be preferred to relative positioning for 3D interfaces, both in general and for 3D modelling in particular. This is in contrast to 2D interfaces where the mouse (relative positioning) is favoured over the tablet (absolute positioning) for many tasks, including much 3D modelling work. The performance of these two interaction techniques for 3D interfaces had not been investigated in previous work and further work is needed to tease out the differences between the 2D and 3D experiences. The study also clearly indicates that head-tracking is not an essential feature for the 3D tasks attempted. This contradicts previous work[4] that suggests a significant effect of head-tracking for similar tasks. Again, further work is needed to confirm whether the results here generalise. If so, a design choice will have to be made since this study showed that, while head-tracking provides no quantitative benefits, it provides a qualitative improvement in the perceived experience.

ACKNOWLEDGMENTS

The authors acknowledge financial support from the EPSRC, grant number EP/E025889/1. Thanks to Ian Davies and Christian Richardt for their help with the equipment.

REFERENCES

[1] Beaton, R. J., DeHoff, R. J., Weiman, N., and Hildebrandt, P. W., "An evaluation of input devices for 3-D computer display workstations," in [*True Three-Dimensional Imaging Techniques and Display Technologies*], *Proc. SPIE* **761**, 94–101 (1987).

[2] Beaton, R. J. and Weiman, N., "User evaluation of cursor-positioning devices for 3-D display workstations," in [*Three-Dimensional Imaging and Remote Sensing Imaging*], *Proc. SPIE* **902**, 53–58 (1988).

[3] Spain, E. H., "Stereo advantage for a peg-in-hole task using a force-feedback manipulator," in [*Stereoscopic Displays and Applications*], *Proc. SPIE* **1256**, 244–254 (1990).

[4] Grossman, T. and Balakrishnan, R., "An evaluation of depth perception on volumetric displays," in [*Proceedings of the Working Conference on Advanced Visual Interfaces*], *AVI '06* , 193–200, ACM, New York, NY, USA (2006).

[5] MacKenzie, I. S., Sellen, A., and Buxton, W. A. S., "A comparison of input devices in element pointing and dragging tasks," *CHI '91* , 161–166, ACM, New York, NY, USA (1991).

[6] Stevenson, D., Smith, K., McLaughlin, J., Gunn, C., Veldkamp, J., and Dixon, M., "Haptic workbench: A multisensory virtual environment," in [*Stereoscopic Displays and Virtual Reality Systems VI*], *Proc. SPIE* **3639**, 356 (1999).

[7] Grant, B., Helser, A., and Taylor, R., "Adding force display to a stereoscopic head-tracked projection display," in [*Proceedings of the Virtual Reality Annual International Symposium*], 81–88, IEEE (1998).

[8] Ikits, M., Hansen, C. D., and Johnson, C. R., "A comprehensive calibration and registration procedure for the Visual Haptic Workbench," in [*Proceedings of the Workshop on Virtual Environments 2003*], *EGVE '03*, 247–254, ACM, New York, NY, USA (2003).

[9] Zhai, S. and Milgram, P., "Human performance evaluation of isometric and elastic rate controllers in a six-degree-of-freedom tracking task," in [*Telemanipulator Technology and Space Telerobotics*], *Proc. SPIE* **2057**, 130–141, SPIE (1993).

[10] Zhai, S., Milgram, P., and Drascic, D., "An evaluation of four 6 degree-of-freedom input techniques," in [*CHI '93: INTERACT '93 and CHI '93 conference companion on human factors in computing systems*], 123–125, ACM, New York, NY, USA (1993).

[11] Zhai, S., "Investigation of feel for 6DOF inputs: Isometric and elastic rate control for manipulation in 3D environments," in [*Computer Systems*], *Human Factors and Ergonomics Society Annual Meeting Proceedings* **37**, 323–327 (1993).

[12] Bowman, D. A., Kruijff, E., LaViola, J. J., and Poupyrev, I., [*3D User Interfaces: Theory and Practice*], Addison-Wesley, Redwood City, CA, USA (2004).

[13] Liang, J., Liang, O., and Green, M., "Geometric modeling using six degrees of freedom input devices," in [*3rd Int'l Conference on CAD and Computer Graphics*], 217–222 (1993).

[14] Nielsen, J., "Usability inspection methods," in [*CHI '94: Conference companion on human factors in computing systems*], 413–414, ACM, New York, NY, USA (1994).

[15] Campbell, D. A. and Stanley, J., [*Experimental and Quasi-Experimental Designs for Research*], Wadsworth Publishing (1963).

[16] Zhai, S., "User performance in relation to 3D input device design," in [*Proc. SIGGRAPH*], *Computer Graphics* **32**(4), 50–54, ACM, New York, NY, USA (1998).