

Relative and Absolute Mappings for Rotating Remote 3D Objects on Multi-Touch Tabletops

Shamus P. Smith
Durham University
Durham, DH1 3LE, UK
shamus.smith@durham.ac.uk

Elizabeth L. Burd
Durham University
Durham, DH1 3LE, UK
liz.burd@durham.ac.uk

Linxiao Ma
Durham University
Durham, DH1 3LE, UK
linxiao.ma@durham.ac.uk

Iyad Alagha
Durham University
Durham, DH1 3LE, UK
iyad.alagha@durham.ac.uk

Andrew Hatch
Durham University
Durham, DH1 3LE, UK
andrew.hatch@durham.ac.uk

The use of human fingers as an object selection and manipulation tool has raised significant challenges when interacting with direct-touch tabletop displays. This is particularly an issue when manipulating remote objects in 3D environments as finger presses can obscure objects at a distance that are rendered very small. Techniques to support remote manipulation either provide absolute mappings between finger presses and object transformation or rely on tools that support relative mappings to selected objects. This paper explores techniques to manipulate remote 3D objects on direct-touch tabletops using absolute and relative mapping modes. A user study was conducted to compare absolute and relative mappings in support of a rotation task. Overall results did not show a statistically significant difference between these two mapping modes on both task completion time and the number of touches. However, the absolute mapping mode was found to be less efficient than the relative mapping mode when rotating a small object. Also participants preferred relative mapping for small objects. Four mapping techniques were then compared for perceived ease of use and learnability. Touchpad, voodoo doll and telescope techniques were found to be comparable for manipulating remote objects in a 3D scene. A flying camera technique was considered too complex and required increased effort by participants. Participants preferred an absolute mapping technique augmented to support small object manipulation, e.g. the voodoo doll technique.

Keywords: 3D object manipulation, direct-touch displays, relative and absolute techniques, user study.

1. INTRODUCTION

Direct-touch interfaces that take human fingers as the input device have a strong appeal to users as they provide an intuitive and compelling interaction experience without using any intermediary devices (Benko et al., 2006). Currently, multi-touch tabletops that are capable of recognising more than one touch input at the same time have become the mainstream of direct-touch tabletops. Users do not have to wait for their turn to interact with the tabletop but rather can work together on the tabletop simultaneously. However, the physical constraints of human fingers have caused specific usability issues to direct-touch interactions. Firstly, the size of an adult's finger may occlude small objects on a tabletop making them difficult to manipulate (Albinsson and Zhai, 2003; Shen et al., 2006). Secondly, compared to mouse cursor use, the human finger is a pointing device with very low

resolution (Albinsson and Zhai, 2003; Ryall et al., 2006). Thus it can be difficult to use the human finger for precise pointing at a target area on a small object. Therefore, applications that work well with mouse as the input device may not be applicable when running on direct-touch tabletops.

In a 3D application, due to the perspective effect, 3D objects are rendered smaller when placed further away from the user. Therefore, objects located or moved to the far end of a 3D scene can be very small. This is less problematic when a mouse or other high-resolution pointing device is employed to manipulate objects because the pointing cursor is small enough to avoid occlusion and supports precise pointing (see Figure 1). However, it can be difficult to manipulate small objects on direct-touch interfaces when fingers are used as the pointing device. As Figure 1 (right)

shows, the visual feedback is lost as a finger obstructs the target object. Without appropriate visual feedback it is impossible for users to conduct fluid manipulation interaction with target objects at a distance. Furthermore, techniques used to manipulate 3D objects on direct-touch tabletops (Hancock et al., 2007; Hancock et al., 2009; Reisman et al., 2009) typically require more than one finger placed on a target object at the same time, for example to rotate an object along the z-axis. Thus the use of multiple fingers can further obscure target objects at a distance.

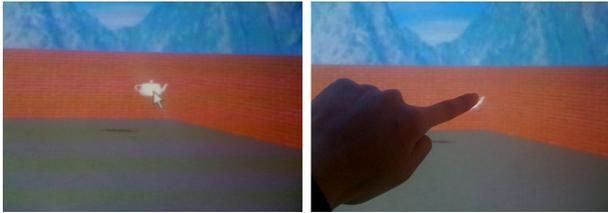


Figure 1: Manipulating a 3D object with a mouse cursor (left) and with a finger (right)

In exploring this problem, we have considered techniques that can fall into two categories, either an *absolute* or a *relative* mapping between interaction technique and the target object. Here, absolute mapping is defined as a co-located, one-to-one correspondence between input position and display position. Relative mapping refers to the correspondence where *motor space*, the space where users' conduct their hand or finger actions, and *display space*, the space that displays visual feedback of an interaction, are not aligned.

Relative mapping techniques have been viewed as an effective way to overcome many of the physical limitations of fingers (Forlines et al., 2006a; Forlines et al., 2006b; Moscovich & Hughes, 2008). The separation of motor space and display space keeps users' fingers and hands away from the target object so that there is no object occlusion. Also a relative mapping mode allows an offset cursor to be mapped to a finger activity. As the offset cursor can have a much higher resolution, it can be used for precise pointing tasks (Albinsson and Zhai 2003; Benko et al. 2006; Potter et al., 1988; Sears & Shneiderman, 1991; Vogel & Baudish, 2007).

While relative mapping has the ability to avoid occlusion and support precise pointing, this mapping mode sacrifices the naturalness of direct-touch interactions. However, with an absolute mapping the same input point on the target object always remains underneath the same fingertip. Therefore, the target object can be manipulated in a more predictable and natural fashion (Sears & Shneiderman, 1991). There is also the concern that an absolute mapping mode would not work

efficiently when the target object is small. To take advantage of the absolute mapping mode the target object has to be enlarged to an appropriate size so that users can still achieve enough visual feedback during the interaction in the context of occlusion from fingers. A number of potential techniques have been proposed in this paper to enlarge target objects.

Two experiments are described to firstly compare the performance of absolute and relative mapping modes for a 3D rotation task. Although previous research (Forlines et al., 2006b; Meyer et al., 1994; Sears & Shneiderman, 1991; Schmidt et al., 2009) has investigated these two mapping modes for 2D manipulation tasks, few studies focus on 3D manipulation tasks. Secondly, a further experiment has explored user preference for four different mapping techniques used to manipulate remote objects in a 3D scene.

2. RELATED WORK

A number of relative mapping techniques have been employed to address the occlusion and low-precision selection issues with direct-touch interfaces that focus on 2D applications. The *Take-Off* technique (Potter et al. 1988; Sears & Shneiderman, 1991) provides users with an offset cursor that is located at a fixed distance above the finger. The selection event is triggered when the user removes their finger from the touch-screen rather than the first contact on the touch-screen. The separation of visual space and motor space ensures the visibility of targets while the use of a high-resolution offset cursor allows the precise selection of small targets. However, Vogel and Baudisch (2007) note that there are drawbacks of this technique. As the touch position is different to the cursor position, users cannot directly aim for the actual target. Also, some display areas of touch-screens, such as the bottom corner, are unreachable for the offset cursor as the offset distance and direction are fixed. To avoid these drawbacks Vogel and Baudisch proposed the *Shift* technique that replaces the offset cursor with an offset callout that displays a copy of the occluded screen area under the finger. With this technique users can directly aim for the target as the touch position is the actual input position and the offset callout only plays the role of an enlarged display space. In addition, although the default position of the callout is above the target it can also be placed elsewhere when the target is close to a screen edge.

Another application of relative mapping is the use of a touchpad that provides users with a separate input space from the display space. A typical example is the touchpad used on laptop computers

that often support single-touch input and plays the role of an external mouse. The touchpads on modern laptops also support multi-touch input that provide users with access to advanced gestures. For example, the touchpad used in Moscovich and Hughes's (2008) study allows users to rotate a 2D object with two fingers.

Another solution to manipulating a small sized object on direct-touch tabletops is to enlarge the target object so that it is suitable for finger-based input. A widely employed method in 2D applications is the zooming mechanism that is often used with a magnifier metaphor. The *Dual Finger Stretch* technique proposed by Benko et al. (2006) allows the non-dominant finger to specify a zooming area that is centered to the dominant finger's location. The dominant finger is then used to perform precise selection actions. This technique has been found to be superior to the other techniques developed by Benko et al. (2006) for precise selection on direct-touch tabletops. Albinsson and Zhai (2003) also conducted a study to evaluate the *ZoomPointing* technique and other non-zooming techniques for precise selection in 2D applications. The results show that ZoomPointing was faster than all other techniques with the same error rate. However, the zooming mechanism with the magnifier metaphor has an unavoidable weakness that the zooming area can occlude the screen area underneath it.

The comparison of relative mapping and absolute mapping is also an active research area but comparative results have been inconclusive. Meyer et al. (1994) conducted a study to compare two absolute mapping input devices, a touch-screen and an absolute pen, and three relative mapping input devices, a mouse, a trackball and a mouse pen. The results indicated that all the relative mapping input devices were faster than the absolute mapping input devices.

Schmidt et al. (2009) compared an absolute mapping input on a tabletop display, where the tabletop is both the input and display device and where the input space is aligned with the display space, with a relative mapping input condition where the tabletop plays the role of input device and a separate screen was used for display. The results show that users performed faster with the absolute mapping input than with the relative mapping input. However, the results also revealed that users can work efficiently with relative mapping input and even inexperienced users can understand the underlying principles of the indirect modality instantly.

Sears and Shneiderman (1991) compared relative mouse input with absolute touch-screen input using a mouse with the control-device gain close to 1 and

a 27.6cm by 19.5cm display. They found that absolute mapping was faster than relative mapping when targets are 16 pixels in width and greater. However, the participants in this study still preferred relative mouse input even when the absolute touch-screen input provided superior performance. Forlines et al. (2006b) compared the absolute and relative pen input on two different display sizes, a Tablet PC and a wall-sized display. They found that absolute input took less selection time on a Tablet PC while relative input outperformed absolute input for distant target selection.

3. TECHNIQUES FOR REMOTE 3D OBJECT MANIPULATION

The manipulation techniques explored in this paper focus on fixed-location 3D manipulations. A number of techniques have been developed for positioning objects in a 3D scene through direct-touch (Hancock et al., 2007; Hancock et al., 2009; Martinet et al., 2009). This paper aims to evaluate several techniques to support such remote manipulation of an object that has been placed at the far end of a 3D scene. The techniques considered either support relative or absolute mapping modes.

3.1 Techniques using relative mapping

The use of relative mapping is a potential way to reach remote 3D objects. A first step could be to generalise the Take-Off (Potter et al., 1988; Sears & Shneiderman 1991) and Shift (Vogel & Baudisch, 2007) techniques for 3D applications. However, while these techniques can support selection tasks in a 2D application, they are less suited for 3D rotation tasks. While a 3D object is displayed on a 2D screen only one side of the object is visible and touchable by user's fingers or a mouse cursor. As it is physically impossible to reach inside of the display, when significantly rotating a 3D object, the user's fingers or mouse cursor will leave the target object. Thus, clutching is often required in order to rotate a 3D object to a target rotation. However, the Take-Off technique cannot directly aim for the actual target. Users have to spend additional effort to aim the cursor on the target object every time they conduct a clutch operation.

The Shift technique does support direct aim for the target object and it is possible to apply this technique to a 3D application by creating and placing a copy of the 3D target object above the finger touch. However, this technique only enlarges the display space. As the motor space is kept unchanged, it can be too small for users to conduct direct manipulation, especially when 3D rotation often involves more than one finger (Hancock et al., 2007; Hancock et al., 2009; Reisman et al., 2009).

An alternative is to provide users with a *virtual touchpad* associated with a target object (see Figure 2). Finger touches and movements on the virtual touchpad are mapped directly to the target object. As the motor space, i.e. the virtual touchpad, is not aligned with the display space, i.e. the actual object, the users' fingers and hands do not occlude the target object. In addition, the motor space provided by the virtual touchpad can be defined to be large enough for users' fingers to perform any required gestures for manipulating the target objects.

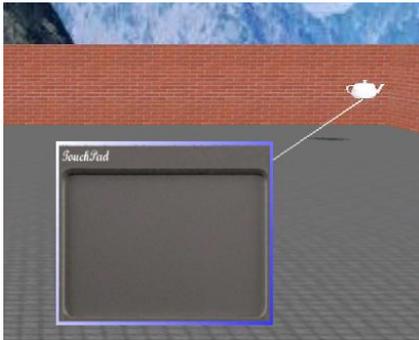


Figure 2: Virtual touchpad

Each touchpad exclusively belongs to one object in the 3D scene and can only be used to manipulate this object. Therefore, different users can use different touchpads to manipulate different objects at the same time. The touchpads can be hidden until the user activates it using an enabling mechanism, for example double clicking on a target object. Also the virtual touchpad can be moved to any position on the screen. This allows users to position an individual touchpad as to not obscure important features of the environment or impact on other users' views in a cooperative environment. Finally, a virtual touchpad can be rotated to any angle in the 2D plane parallel to the tabletop screen. This allows users to use the touchpad on any side of a tabletop device.

One disadvantage of the touchpad technique is that only the motor space is enlarged while the size of the visual space is unchanged. The touchpad would be not suitable to manipulate an object located very far away as this object may be rendered so small that users cannot view it or select it to enable manipulation.

Any gestures for manipulating 3D objects do not have to be hard bound to the touchpad. The techniques explored in this paper are expected to be independent from explicit gestures for object manipulation. Thus developers can design and bind different gestures to a touchpad. Currently, a number of gestures have been suggested for manipulating 3D object on direct-touch interfaces

(Cardinaels et al., 2008; Edelmann et al., 2009; Hancock et al., 2007; Hancock et al., 2009; Reisman et al., 2009). These gestures could, potentially be bound to a touchpad.

3.2 Techniques using absolute mapping

A predominate feature of absolute mapping techniques is that they work on a target, or proxy, object that has a suitable size. A number of techniques for enlarging, representations of, remote 3D objects are discussed in this section.

There are numerous methods for enlarging remote objects in a 3D scene. However, many are not suitable for multi-touch environments, and particularly multi-user multi-touch environments. For example, a technique might push the camera view closer to a target object, e.g. by zooming in or moving the camera towards the object as in the *Head-Butt Zoom* technique (Mine et al., 1997), and then pulling the camera view back to the original position after completing the required manipulation. Unfortunately, such a technique has a high use overhead as users have to frequently adjust the camera view. More importantly, this technique limits environment use to a single user. It is common for tabletop environments to be shared and thus the change of the camera view would affect other users.

Another similar technique to enlarge the target object is to move the object closer to the users. Using 3D perspective, the object would thus be rendered larger. However, this technique involves extra steps to translate the object from and back to the original location. In addition, users may have to fine-tune the rotation of the target object when it is translated back to the original location as the view angle, when rendered closer, is different than when at its original, remote, location. Finally, a technique may support users to directly scale up an object at the original location for direct manipulation. However, any enlarged object may intersect with nearby objects in the scene and cause object occlusion resulting in potential selection errors.

This paper explores three techniques that are suitable for enlarging remote objects on direct-touch tabletops avoiding the issues defined above, namely the *voodoo doll*, *telescope* and *flying camera* techniques.

Similar to Pierce et al.'s (1999) *Voodoo Doll* technique, the *voodoo doll* used here (see Figure 3) is a copy of a target object with the same properties. The only difference is that the "doll" is placed closer to the camera so that it is rendered larger. The doll is not only big enough for the users' fingers to conduct the required gestures but also large enough that the user's fingers do not totally

occlude visual feedback during manipulation. The doll plays the role of a proxy for the target object. Users directly manipulate the doll and the manipulation is synchronously applied to the target object. As with the touchpad, the doll can be moved to any position so that users can work with it at a preferred position without affecting other users.

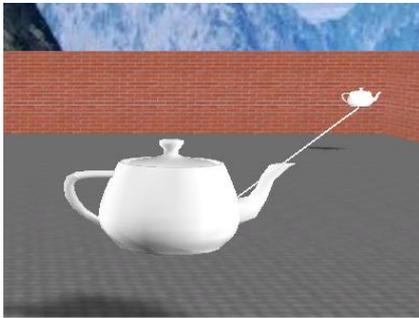


Figure 3: Voodoo doll technique

A telescope is a common device to view distant objects. This paper borrows this metaphor to enlarge remote objects on direct-touch tabletops. A *virtual telescope* technique (see Figure 4) consists of two main components, namely the frame and the lens. Users can place one finger on the frame and drag the virtual telescope to any location on the screen. Objects underneath the virtual telescope are shown on the lens with an enlarged size. Users can directly manipulate the enlarged object through the lens. With a traditional physical telescope, the view can be zoomed in and out in order to achieve an optimal view. A zooming mechanism has also been implemented in the virtual telescope technique. Users can place two fingers on the frame and rotate the frame clockwise to zoom in and rotate anticlockwise to zoom out. Unlike the touchpad and the voodoo doll, the virtual telescope does not exclusively associate with a specific object but rather it can be moved over and enlarge any object in the 3D scene. Also more than one virtual telescope can be activated and used simultaneously on the same tabletop.



Figure 4: Virtual telescope technique

However, the virtual telescope technique has similar limitations to the magnifier metaphor. When the virtual telescope is placed over an object, the

area underneath the telescope is occluded and only a part of this area can be seen through the telescope lens.

A *virtual flying camera* can be implemented on direct-touch tabletops that provide users with an additional viewport into the 3D scene (see Figure 5). Virtual cameras can be moved to any position in the 3D space and rotated to any angle in order to point the lens at different areas of a scene. The view captured by the virtual camera is displayed on a monitor panel. This monitor panel can also be moved to any location in order to avoid the occlusion of other objects. The monitor panel can also be rotated to any angle in the screen plane so that users can use it on any side of a tabletop device. When the virtual camera points at a target object, a copy of that object is shown on the monitor panel. Users can use the control buttons on the monitor panel to zoom in or out (+ and – in Figure 5) in order to adjust the size of the object. Thus users can directly manipulate the proxy object on the monitor panel and the manipulation is immediately applied to the remote target object.



Figure 5: Flying camera technique

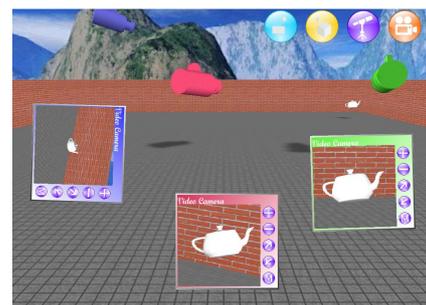


Figure 6: Multiple flying cameras on the same tabletop

Any number of virtual cameras can be used on one tabletop device at the same time (see Figure 6). In addition to supporting remote object manipulation, the virtual camera technique can address another important issue with direct-touch tabletops (Grossman & Wigdor, 2007). While displaying a 3D object on a 2D screen, users can only see one side of the object. However, users often want to observe the object from different sides or angles, especially

when a group of users are sharing the same 3D scene and seated around the same tabletop device. The virtual cameras solve this problem as they can be positioned and pointed at any place in the 3D space. However, the capability of virtual cameras to support multiple views is out of the scope of this paper and we focus here on the virtual cameras capability for manipulating remote objects.

4. ABSOLUTE VERSUS RELATIVE MAPPING MODES FOR A 3D ROTATION TASK

An experiment (E1) was conducted to compare the performance of absolute mapping and relative mapping modes for a 3D rotation task with different sized objects.

4.1. Method

Twenty volunteers (16 male and 4 female, age range 18-52) were recruited from the local university. All participants were right hand dominant. No participant had frequent experience of direct-touch tabletops. Ten participants had daily or weekly usage of a touch-based mobile phone. Only one participant reported frequent experience of 3D applications.

The experiment was conducted on a rear-projected multi-touch tabletop that was set at 1024*768 pixels resolution with a 86cm*64cm display area which yields a pixel size of 0.84mm. A finger that is about 1.5cm wide covers about 18 screen pixels. Diffused illumination principle was employed for the touch detection using a Point Grey Firefly MV camera with a resolution of 640*480 at 60fps to capture points of touch. The height of the table was 88cm and the screen was tilted to 30°. Participants stood in front of the tabletop.

Two independent variables were investigated, including mapping mode (absolute mapping and relative mapping) and size of the target object (6.5cm in width, 13cm in width and 19.5cm in width). Even the smallest size is big enough for users (with a normal finger size) to conduct the required manipulation with enough visual feedback through the fingers. Sears and Shneiderman (1991) found that target size can impact the performance of absolute mapping and relative mapping on 2D interaction. This experiment investigated this impact on 3D interaction.

The participants were asked to rotate a 3D teapot from the start rotation (135° about the x-y-z-axis) to the target rotation (0° about the x-axis, y-axis, and z-axis). When the teapot was rotated to the target rotation within the given tolerance, the colour of the teapot changed to silver and the teapot was frozen. In the absolute mapping condition, participants

were asked to directly manipulate the teapot where the first contact point on the target object always remains underneath the same fingertip (see left in Figure 7). In the relative mapping condition, participants were asked to use the virtual touchpad to indirectly manipulate the teapot (see right in Figure 7). The control-device gain of the touchpad was set to 1 for all sizes so that the same finger movement distance on the touchpad and the target object would cause the same angle change on the object. Both conditions employed the same gesture to manipulate the target object, i.e. using one finger to rotate the object about x-axis and y-axis, and using two fingers for rotating the object about the z-axis.



Figure 7: The experimental setting for absolute mapping condition (left) and relative mapping condition (right)

A within-subject design was used. There were 5 trials for each condition. A total of 2 (mapping mode) X 3 (target size) X 5 (repetitions) = 30 test trials were collected from each participant. The order of conditions was counterbalanced across the participants. The task completion time, the number of touches and participant preference were collected for analysis.

4.2 Results

There was little difference between the average task completion times and the number of touches across all object sizes (see Figure 8). Although the average relative mapping time was slightly less and the number of relative mapping touches slightly more, this was not statistically significant.

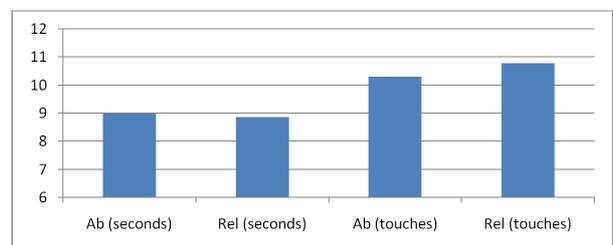


Figure 8: Average times and touches over all object sizes with the absolute and relative mapping modes

As seen in Figures 9 and 10, when participants used the relative mapping mode the task completion time increased as the size of the target object increased. This is expected as the

participants' finger had to travel further on the tabletop surface to rotate larger objects to a given rotated image compared to rotating smaller objects. This linear relation was also present for the number of touches with the relative mapping condition. However, neither condition was statistically significant.

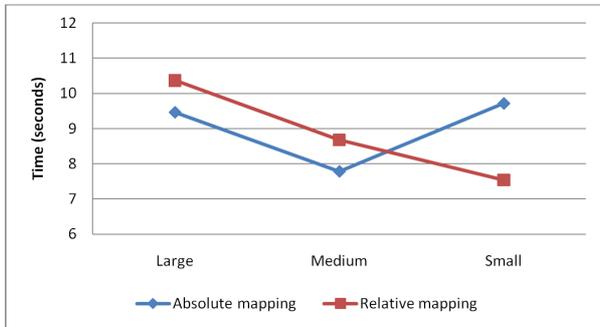


Figure 9: Average times over large, medium and small objects with the absolute and relative modes

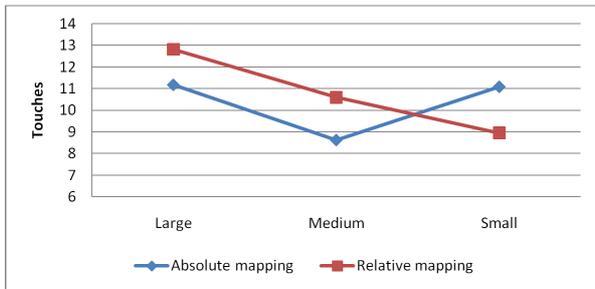


Figure 10: Average touches over large, medium and small objects with the absolute and relative modes

A similar relation is seen with the absolute mapping with the exception of small objects. This could be caused by the participants' fingers obscuring the visual feedback and thus hindering the task completion. To investigate this difference, the data from the medium and large size conditions was merged and compared to the small size condition. This showed an interaction affect between mapping modes and target size that was significant for both completion time ($F=5.383$, $p=0.022$) and number of touches ($F=5.463$, $p=0.021$). Therefore, in the context of this experiment, absolute mapping required less time and fewer touches for medium and large objects but not for small objects.

The participants were also asked about their mode preference after they had completed all the trials. Although the absolute mapping can provide users with more "natural" interaction experience, overall there was little difference in preference (relative mapping 47%; absolute mapping 53%). However, within object size difference, participants had a clear preference for the relative mapping mode with small objects (70%) and the absolute mapping mode for medium-sized objects (70%). When large

objects were to be manipulated the preference difference was reduced (relative mapping 40%; absolute mapping 60%). Thus participants only had a preference for the absolute mapping mode with medium and large objects.

5. COMPARING MAPPING TECHNIQUES FOR REMOTE 3D OBJECT MANIPULATION

A second experiment (E2) was conducted to compare the usability of different mapping techniques. Due to the perspective effect of 3D virtual scenes, we were particularly interested in the manipulation of remote objects that would be rendered small at a distance.

5.1 Method

The same participants and equipment used for experiment E1 (Section 4.1) were used.

Participants were required to rotate a remote 3D object, a teapot, in a 3D virtual scene to a set rotation using the techniques described in Section 3, namely touchpad, voodoo doll, telescope and flying camera techniques.

When using the touchpad to rotate the target object, participants were given the freedom to place the touchpad at any location on the screen. As the touchpad is a 2D object that is rendered in orthogonal mode, its view angle is the same no matter where it is been placed. In comparison, the doll in the voodoo doll technique is itself a 3D object. Thus the view angle of the doll can be changed when it is moved to other positions. Therefore, the participants were restricted to using the doll representation at two set locations, namely the center of the table where the participants saw a straight-front view of the doll and the left end of the table where the participants saw a right-front view of the doll.

When the virtual telescope technique was activated, it was initially located at the bottom-left corner of the scene, while the target object was placed close to the top-right corner. Participants were asked to position the telescope over the target object and adjust the view of the telescope in order to achieve a suitable size for the target object by zooming in or out. The telescope view was initially set at an extreme distance to the target object thus forcing the participants to actively use the zoom mechanism in the lens before they could manipulate the remote object.

For the flying camera technique, participants were asked to translate the camera to three positions in the 3D scene and point the camera at the target object. The first position was at the front of the target object. This position was in the same XY

plane with the camera so the participants did not have to move the camera along the z-axis. The second position was at the left side of the 3D scene. Also the camera was only required to be moved in the same XY plane. The final position was at the far left side of the 3D scene. The participants not only needed to move the camera in the XY plane but also needed to move the camera along the z-axis.

The gestures used to rotate the target object were the same as those used in experiment E1, i.e. one finger movement caused the object rotation about x-axis and y-axis and two finger rotation caused the object rotation about z-axis.

Two Semantic Differential Scale (Preece et al., 2007, pg314) questions were used to investigate how easy the participants felt it was to use and learn the techniques. The rating can be converted to a number scale: from 1 (difficult to use/learn) to 9 (easy to use/learn). Also the participants were asked about their preference of the four techniques.

5.2 Results

All the participants successfully completed the rotation tasks using all the techniques. Perceived ease of use and learnability results can be seen in Figures 11 and 12 respectively.

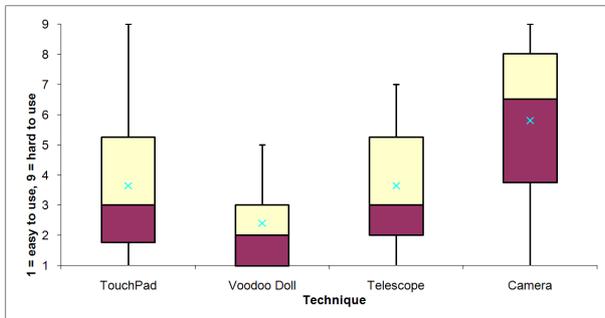


Figure 11: Mapping technique by ease of use results

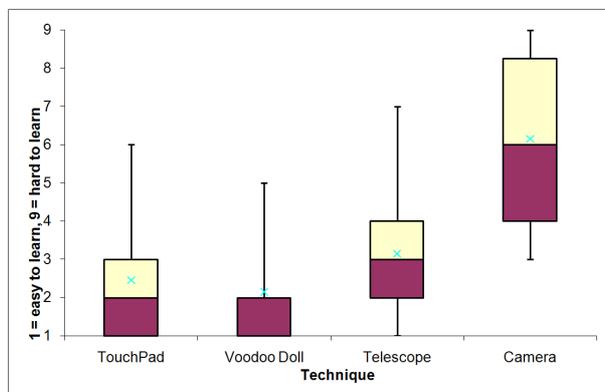


Figure 12: Mapping technique by learnability results

The flying camera technique was found to be the worst performing technique in both categories. Most participants indicated that this technique was difficult to use (Mean = 5.8, SD = 2.67) and difficult to learn (Mean = 6.15, SD = 2.25)¹, and only one participant preferred to use it. Comments from the participants reveal that they had to spend more effort to use the flying camera compared to the other techniques. To manipulate the target object, participants had to set the camera location and the camera view before they could start to manipulate the target object via the monitor screen. The participants considered that these extra steps made the flying camera more difficult to use and learn.

In addition, the participants had to frequently switch their attention between the main viewport of the 3D scene and the viewport of the flying camera in order to consider the view angles of the target object in the two viewports. The view angle of the target object in the main viewport can be very different than the view angle in the viewport of the flying camera (see Figure 13), especially when the flying camera was not placed at a straight angle to the target object. This required participants to constantly switch their attention between the viewports.

The touchpad was identified as easy to use (M = 3.65, SD = 2.41) and easy to learn (M = 2.45, SD = 1.61). Similar to the results observed from experiment E1, the separation of visual space and motor space did not divert the participants' attention. No participants reported difficulty in mirroring their hand manipulation on the touchpad to the visual feedback they observed on the target object. However, participants did comment that the remote teapot was too small to observe as the touchpad did not enlarge the visual space of the target object.



Figure 13: Difference of object view angles between the viewports

¹ As the rating scales are from 1 to 9, this paper defines techniques as difficult to use/learn when the mean > 4.5 and defines techniques as easy to use/learn when the mean <= 4.5.

The voodoo doll was the overall best performer. This technique achieved the lowest values for ease of use (Mean = 2.4, SD = 1.35) and learnability (Mean = 2.15, SD = 1.31). Participants commented that the voodoo doll technique was straightforward to use as the object to be manipulated was just an enlarged copy of the target object. In addition, participants showed no difficulties in using the voodoo doll technique at both the pre-defined positions (e.g. center and left-end of the tabletop). The participants did not comment on any problems with the difference of view angle when rotating the target object to the required rotation.

Although the use of the telescope technique involved extra steps for positioning and zooming, these steps were found to be more straightforward when compared to the positioning and zooming operations required for the flying camera technique. Similar to the touchpad and voodoo doll techniques, this technique was marked as easy to use (Mean = 3.65, SD = 1.95) and easy to learn (Mean = 3.15, SD = 1.76). However, participants did note some usability issues. Firstly, as the target object was manipulated through the telescope lens, some parts of the target object could be rotated out of the current lens view. Thus participants had to reposition the telescope in order to make final adjustments to the teapot at centre. Secondly, one participant accidentally moved the telescope when wanting to manipulate the object (e.g. one finger accidentally dragged across the telescope frame). The participant suggested that the telescope should be locked when the target object is being manipulated.

The overall participant preference resulted in the voodoo doll technique as the favourite technique, with 60% preference, and the flying camera, at 5%, as the most disliked. The telescope and touchpad, at 20% and 15% preference respectively, had similar results. Thus the participants in this experiment have indicated a preference for simple over complex techniques and direct mappings between target and manipulated objects when using mapping techniques.

6. CONCLUSIONS AND FUTURE WORK

This paper has described two experiments to explore techniques to manipulate remote 3D objects on direct-touch tabletops. Both absolute and relative mapping modes were considered. The experiments found that the main effects of these two mapping modes were not statistically significant on both task completion time and the number of touches. This result implies that the separation of motor space and visual space of target objects does not affect users' performance on 3D rotation tasks. Users have the ability to map

their hand actions in the motor space to the visual feedback in the visual space. A similar example is the use of a mouse device where users do not pay visual attention to their hand but rather can focus on the cursor movements on screen. However, the relative mapping mode was found to be faster than the absolute mapping mode when rotating small objects. Also, more participants preferred the relative mapping mode with small target objects. This result is not surprising as users may obscure small objects with their fingers in the absolute mapping mode. However, when the size of target object is medium or large, performance is similar over both mapping methods.

Four different mapping mode techniques were then compared. With the exception of the flying camera technique, the other techniques, i.e. touchpad, voodoo doll, and telescope, were found to be easy to use and learn when manipulating a remote object in a deep 3D scene. Although the touchpad technique separates the motor space of the target object from its visual space, the participants' attention was not diverted. One weakness of the touchpad technique was that it can only enlarge the motor space of the target object. When the target object is placed very far away, the touchpad is not preferable as the visual space of the target object is so small that user cannot view the details of the target object.

The voodoo doll was found to be straightforward to use and the preferred mapping technique. Compared to the touchpad, the voodoo doll enlarges both the motor space and the visual space. Although the telescope and flying camera techniques also enlarge both the motor space and the visual space, they require extra operations for positioning and view setting. When these operations are present, users preferred the simplicity of the telescope technique when compared to the added complexity of the flying camera technique. With the flying camera technique, any inconsistency of the main viewport and the camera viewport caused additional work for the users. However, the flying camera can support interaction in specific scenarios where the other techniques considered here would not be suitable, for example, when the target object is occluded by other objects in the 3D scene or the user requires one or more additional views of the same target object from different angles. How this would affect the usability preference of users is the focus of ongoing work. Also the work explored the perceived usability of mapping techniques for a single-user task. Given the cooperative nature of many tasks on multi-touch tables, future work will investigate how users might use the techniques described here in a shared multi-user context.

7. ACKNOWLEDGEMENTS

The authors would like to thank the experiment participants. This work was partially funded by the EPSRC ERSC TLRP **SynergyNet** project (Grant RES-139-25-0400).

REFERENCES

- Albinsson, P.-A. and Zhai, S. (2003) High precision touch screen interaction. *ACM CHI Conference on Human Factors in Computing Systems*, 105-112.
- Benko, H., Wilson, A. D., and Baudisch, P. (2006) Precise selection techniques for multi-touch screens. *ACM CHI Conference on Human Factors in Computing Systems*, 1263-1272.
- Cardinaels, M., Frederix, K., Nulens, J., Van Rijsselbergen, D., Verwaest, M., and Bekaert, P. (2008) A multi-touch 3D set modeler for drama production. *Proceedings of the International Broadcasting Convention Conference*, 330-335.
- Edelmann, J., Schilling, A., and Fleck, S. (2009) The DabR - A multitouch system for intuitive 3D scene navigation. *3DTV Conference: The True Vision - Capture, Transmission and Display of 3D Video*, 1-4.
- Forlines, C., Vogel, D., and Balakrishnan, R. (2006a) HybridPointing: fluid switching between absolute and relative pointing with a direct input device. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*, 211-220.
- Forlines, C., Vogel, D., Kong, N. and Balakrishnan, R. (2006b) Absolute vs. relative direct pen input. *Mitsubishi Electric Research Labs Tech Report*, TR2006-066.
- Grossman, T., and Wigdor, D. (2007) Going deeper: a taxonomy of 3D on the tabletop. *Proceedings of the Second IEEE International Workshop on Horizontal Interactive Human-Computer Systems*, 137-144.
- Hancock, H., Carpendale, S., and Cockburn, A. (2007) Shallow-depth 3D interaction: Design and evaluation of one-, two- and three-touch techniques. *ACM CHI Conference on Human Factors in Computing Systems*, 1147-1156.
- Hancock, H., Cate ten, T., and Carpendale, S. (2009) Sticky tools: full 6DOF force-based interaction for multi-touch tables. *ACM Interactive Tabletops and Surfaces*, 145-152.
- Hancock, H., Hilliges, O., Collins, C., Baur, D., and Carpendale S. (2009) Exploring tangible and direct touch interfaces for manipulating 2D and 3D information on a digital table. *ACM Interactive Tabletops and Surfaces*, 77-84.
- Martinet A., Casiez G., and Grisoni L. (2009) The design and evaluation of 3D positioning techniques for multi-touch displays. *INRIA Research report*. Number 7015, INRIA.
- Meyer, S., Cohen, O. and Nilsen, E. (1994) Device comparisons for goal-directed drawing tasks. *CHI'94 extended abstracts*, 251-252.
- Moscovich, T. and Hughes, J. F. (2008) Indirect mappings of multi-touch input using one and two hands. *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems*, 1275-1284.
- Mine, M. R., Brooks Jr, F. P. and Sequin, C. H. (1997) Moving objects in space: Exploiting proprioception in virtual-environment interaction. *ACM SIGGRAPH 97*, 19-26.
- Pierce, J. S., Stearns, B. C. and Pausch, R. (1999) Voodoo dolls: Seamless interaction at multiple scales in virtual environments, *I3D '99 Proceedings of the 1999 symposium on Interactive 3D graphics*, 141-145.
- Preece, J., Rogers, Y., and Sharp, H. (2007) *Interaction Design: Beyond Human-Computer Interaction*, Second edition, Wiley: New York.
- Potter, R.L., Weldon, L.J., and Shneiderman, B. (1988) Improving the accuracy of touch screens: An experimental evaluation of three strategies. *ACM CHI Conference on Human Factors in Computing Systems*, 27-32.
- Reisman, J. L., Davidson, P. L., and Han, J. Y. (2009) A screen-space formulation for 2D and 3D direct manipulation. *Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology*, 69-78.
- Ryall, K., Forlines, C., Shen, C., Morris, M., R., and Everitt, K. (2006) Experiences with and observations of direct-touch tabletops. *First IEEE International Workshop on Horizontal Interactive Human-Computer Systems (Tabletop'06)*, 89-96.
- Sears, A. and Shneiderman, B. (1991) High precision touchscreens: Design strategies and comparisons with a mouse. *International Journal of Man-Machine Studies* **34**:4, 593-613.
- Schmidt, D., Block, F., and Gellersen, H. (2009) A comparison of direct and indirect multi-touch input for large surfaces. *INTERACT'09 Proceedings of the 12th IFIP TC 13 international conference on Human-Computer Interaction: Part I*, 582-594.
- Shen, C., Ryall, K., Forlines, C., Esenther, A., Vernier, F. D., Everitt, K., Wu, M., Wigdor, D., Morris, M. R., Hancock, M., and Tse, E. (2006) Informing the design of direct-touch tabletops. *IEEE Computer Graphics and Applications* **26**:5, 36-46.
- Vogel, D. and Baudisch, P. (2007) Shift: a technique for operating pen-based interfaces using touch. *CHI '07 Proceedings of the SIGCHI conference on Human factors in computing systems*, 657-666.