Journal of Universal Computer Science, vol. 24, no. 9 (2018), 1217-1243
submitted: 30/12/17, accepted: 30/6/18, appeared: 28/9/18 © J.UCS

Target Selection in Head-Mounted Display Virtual Reality Environments

Difeng Yu
(Department of Computer Science and Software Engineering
Xi'an Jiaotong-Liverpool University, Suzhou, China
Difeng.Yu14@student.xjtlu.edu.cn)

Hai-Ning Liang*
(Department of Computer Science and Software Engineering and AI University Research Centre (AI-URC)
Xi'an Jiaotong-Liverpool University, Suzhou, China
HaiNing.Liang@xjtlu.edu.cn
*corresponding author)

Feiyu Lu
(Department of Computer Science and Software Engineering
Xi'an Jiaotong-Liverpool University, Suzhou, China
Feiyu.Lu14@student.xjtlu.edu.cn)

Vijayakumar Nanjappan
(Department of Computer Science and Software Engineering
Xi'an Jiaotong-Liverpool University, Suzhou, China
Vijayakumar.N@xjtlu.edu.cn)

Konstantinos Papangelis
(Department of Computer Science and Software Engineering
Xi'an Jiaotong-Liverpool University, Suzhou, China
K.papangelis@xjtlu.edu.cn)

Wei Wang
(Department of Computer Science and Software Engineering
Xi'an Jiaotong-Liverpool University, Suzhou, China
Wei.Wang03@xjtlu.edu.cn)

Abstract: Target selection is one of the most common and important tasks in interactive systems. Within virtual reality environments, target selection can pose extra challenges to users because targets can be located far away, clustered together, and occluded from view. Although selection techniques have been explored, it is often unclear which techniques perform better across different environmental target density levels and which have higher levels of usability especially for recently released commercial head-mounted display (HMD) virtual reality systems and input devices. In this paper, we first review previous studies on target selection in HMD VR environments. We then compare the performances of three main techniques or metaphors (RayCasting, Virtual Hand, and Hand-Extension) using recently marketed VR headsets and input devices under different density conditions and selection areas. After, we select the best two techniques (RayCasting and Virtual Hand) for the second experiment to explore their relative performance and usability by adding different feedback to these two

techniques. In the third experiment, we implemented three techniques with pointing facilitators and compared them against the best techniques from the second experiment, RayCasting with visual feedback, to assess their performance, error rates, learning effects, and usability. The three studies altogether suggest the best target selection features, based on techniques, feedback, and pointing facilitators for target density conditions in HMD VR environments.

Keywords: Virtual Reality; 3D User Interaction; Target Selection; Occlusion; Oculus RIFT; HTC Vive
Categories: D.0, B.4.2, J.4, L.3.1

1 Introduction and Background

Recent years have witnessed the proliferation of virtual reality technologies marketed to the masses. In parallel, there are new motion tracking and input devices that enable richer spatial interactions within VR environments [Argelaguet, 03] [Zaranet, 14] [Farmani, 17] [Mayer, 18]. Techniques with different features provide users with multiple choices to interact with the objects in virtual environments [Liang, 16]. One fundamental interaction is target selection, the initial task for most common user activities [Kulik, 09]. Target selection can be seen in a wide range of applications. With advancements in VR, target selection can become non-trivial. Applications in gaming, for example, require users to select targets quickly and accurately. In addition, with higher resolution head-mounted displays (HMD), a greater number of targets can fit into smaller areas [Liang, 18]. Occlusion of targets can increase the difficulty of quick and precise selection. In this research, we aim to explore target selection in virtual reality environments that are becoming popular with consumer devices such as Oculus RIFT and HTC Vive.

In the context of target selection, three aspects are important: (1) selection metaphors; (2) feedback; and (3) pointing facilitators. Selection metaphors can be divided broadly into 2 categories: virtual hand [Mine, 95] and virtual pointing [Fitzmaurice, 93] [Liang, 94]. Virtual hand techniques map the real hand or hand-held devices as the virtual cursor in 3D coordinates. This results in a more natural interaction with objects in a virtual environment. With the introduction of the virtual arm [Poupyrev, 96], virtual hand techniques overcome the constraints of the real world by enlarging the selection area. In addition to virtual hand techniques, virtual pointing techniques also allow users to select the objects beyond their reach and require relatively less physical movement [Bowman, 97] [Argelaguet, 03]. RayCasting is considered as one of the most popular techniques in 3D selection tasks. It works by casting a virtual ray into the 3D environment and makes the ray interact with the target of interest [Vanacken, 07]. This technique is commonly seen in VR headsets or hand-held devices. However, these techniques may suffer from the object occlusion when the environment has a high density of targets [Stoakley, 95].

Visual, haptic, and audio feedback are usually used in conjunction with selection techniques to guide users with their selection task and have shown to be useful [Akamatsu, 95]. Also as pointed out by Wingrave et al. [Wingrave, 02], feedback is an essential part of virtual environments and users may not be able to interact with the objects efficiently without it. However, simply including feedback may not necessarily lead to positive user experiences [Wingrave, 05].

Pointing facilitators are enhancements for efficient selection of targets. Following the Fitts' law [Fitts, 54] [MacKenzie, 92], as an object's size and location have a direct effect on selection performance, several techniques have focused on increasing the object size or decreasing the amplitude of movement [Liang, 94] [Forsberg, 96] [Vanacken, 07]. Most of these pointing facilitators have been shown to improve the selection performance in the space where the targets are fairly sparsely distributed. However, when dealing with high-density targets, their effectiveness is less certain. False positive activation of unwanted targets negatively affects efficiency and thus user experience.

This research focuses on what techniques, feedback types, and pointing enhancements will work best for users to select targets in HMD VR environments in both sparse and dense environments. Some prior research has explored different techniques' performance in the immersive 3D virtual environment (e.g., see [Bowman, 97] [Cournia, 03] [Teather, 11]). However, only a few have looked at the performance of techniques in different levels of target density (i.e., the number of distractors surrounding the goal target within a certain range), target overlapping or occlusion, and the learning effect of the techniques. As the proliferation of VR HMD continues, it is important to examine how we can support the effective selection of objects within these environments.

The main aim of this research is three-fold:

1. To determine which technique(s) and features that will allow for the fastest selection with relatively low selection misses in simple environments (low-density level);
2. To determine which technique(s) and features that will allow for the fastest selection with relatively low selection misses in complex environments (typically with high-density with occluded targets);
3. To propose recommendations for designers to design selection techniques for recently marketed HMD VR systems.

In the remainder of this paper, we first review previous research in target selection. Then we present our motivation for the work and three studies that are built on each other. In the first study, we explore the performance of three pointing metaphors or techniques (RayCasting, Virtual Hand, and Hand-Extension) in different density conditions and selection areas. In the second study, we select the two best techniques, which are RayCasting and Virtual Hand, and explore the addition of feedback to these techniques to determine if feedback can increase their performance and lower error rates. In the third study, we compare the best performing technique with added feedback against techniques that are enhanced with pointing facilitators to see their comparative selection speed, error rate, and learning effect. At the end of the paper, we summarize the findings of the three experiments and provide distilled lessons.

2 Related Work

Selection (along with movement, manipulation, and scaling) is categorized as one of the four elementary interaction tasks not only in hybrid and mixed reality spaces

[Papangelis, 17] but also in a virtual environment [Mine, 95]. Object selection allows the user to acquire or identify a target for future interaction so that other interactions can be performed on or with it. Because it is such an important task, there have been studies on various aspects of target selection [Vanacken, 06] [Cockburn, 11] [Lubos, 14]. In this section, we will first review some main selection metaphors for 3D environments. After that, we will provide an overview of the literature related to feedback types. Lastly, we will describe pointing facilitators that aim at shortening the object selection times.

2.1 3D Selection Metaphors

In the context of task selection, there are two main egocentric metaphors: (1) virtual hand [Lee, 16] [Tran, 17] and (2) virtual pointer [Poupyrev, 98] [Poupyrev, 99].

2.1.1 Virtual Hand. With virtual hand, users are able to select the target objects by first "touching" them with a virtual representation of their real hands and then use a trigger for selection (such as pressing a button, issuing a voice command, or making a hand gesture) to confirm the selection [Bowman, 04]. This makes interaction simple by mapping virtual tasks with real tasks which results in a more natural human interaction. The classical implementation of this metaphor is the direct mapping of the user's hand movement to the virtual hand's motion in a virtual environment, while other techniques, such as the "Go-Go" [Poupyrev, 96], uses non-linear mapping and breaks away from physical constraints to give greater range of actions—the Go-Go technique, for example, grows the user's arm to allow greater reach.

2.1.2 Virtual Pointer. Unlike the virtual hand metaphor, virtual pointer allows selection of objects by pointing at them—e.g., seen as a virtual ray in the environment. Users can pick different objects by manipulating the start point and orientation of the ray. The ray is estimated from the orientation and the position of the user's virtual hand. Alternatively, it may emanate from the tracked head position and extended in the direction to which the head is pointing (known as "Head-Based RayCasting") [Bowman, 04] [Qian, 17] [Kytö, 18]. RayCasting [Mine, 95] is a common implementation of the pointer metaphor. Some more sophisticated implementations include Aperture [Forsberg, 96], Flashlight [Liang, 94], Shadow Cone Selection [Steed, 04], and the techniques for cluttered virtual environments [Argelaguet, 09].

There has been some work that has compared the performance of RayCasting techniques and Virtual Hand ones in immersive environments [Poupyrev, 98]. This research suggests that RayCasting could be more efficient when selecting large-size and close-distance objects but it may not perform well when the high precision of selection is required [Poupyrev, 98] [Bowman, 01]. On the contrary, Virtual Hand techniques can achieve comparable performance for all conditions of local manipulation (within the area of the user's maximum reach) [Poupyrev, 98]. However, this work does not take into account the number of errors caused by the two selection techniques during the selection task and is less focused on comparing selection metaphors for different selection with density levels of targets in VR according.

2.2 Feedback

Existing literature posits that providing feedback may be beneficial to selection tasks in 3D virtual environments [Herndon, 94] [Vanacken, 06] [Teather, 14] [Ebrahimi,

16]. As outlined by Bowman et al. [Bowman, 01], feedback can be divided into three categories: graphical feedback, force/tactile/haptic feedback, and audio feedback.

2.2.1 Visual Feedback

There are two common forms of visual feedback [Grossman, 06]. The first is target highlighting [Mine, 95] [Mould, 04] which highlights the object or its bounding box, and the other one is called target shadowing [Wanger, 92] which projects both target and cursor on a ground plane. This enhancement allows the user to make sure they select the right target, thus improving the accuracy of selection. However, previous research has found that it does not always improve user performance [Poupyrev, 98] [Wingrave, 05] [Guillon, 16].

2.2.2 Haptic Feedback

Providing haptic feedback such as vibration and bump can also help users improve selection accuracy [Vanacken, 06] [Pavlovych, 09] [Pfeiffer, 15]. Research in 3D immersive environments has indicated that the addition of haptic feedback can improve tapping performance and also reduce errors [Arsenault, 00]. On the other hand, in dense environments where distractor targets are introduced, the haptic feedback may become counterproductive as the user might be guided to the wrong object [Wanger, 92].

2.2.3 Audio Feedback

Vanacken et al. [Vanacken, 06] have found that audio feedback can improve the selection reaction time in 3D environments. For example, it can notify the users when a target is under selection by playing an audio earcon sound for a short duration. However, similar to haptic feedback, auditory feedback may become annoying and impractical in dense environments.

There has been only very little work that has explored the effect of feedback on 3D selection tasks [Vanacken, 06]. Vanacken et al. compared force, audio, visual feedback and showed that only visual feedback showed to be useful. However, their experiment was done in non-immersive LCD shutter glasses; and the results cannot be assumed to hold in immersive virtual reality. There also has been scant research which has discussed the comparative performance of different types of feedback to both two main metaphors (which is pointing and virtual hand). The combination of different types of feedback has not been sufficiently explored in this area.

2.3 Points Facilitators

Following Fitts' law [Fitts, 54] [MacKenzie, 92], several techniques have been proposed to increase user's performance in selection tasks. The usual form of Fitts' law is formulated as:

$$\mathrm{MT} = \mathrm{a} + \mathrm{b}log_2(\frac{A}{W} + 1) \quad (1)$$

It states that the movement time (MT) to select a target can be calculated with

width W and distance A from the cursor [MacKenzie, 92]. The 'a' and 'b' are constants and the logarithmic term is called index of difficulty (ID) of a motor task. To decrease the movement time, facilitators attempt to either increase the object size (W) or decrease the amplitude of movement (A).

One of the most successful pointing facilitators is the Bubble Cursor [Grossman, 05] which enables target selection by dynamically resizing of the cursor's activation area and significantly outperforms the traditional pointing cursors. This technique divides the activation boundaries into Voronoi cells and the objects can be selected when the selection tool is inside the Voronoi cell which encloses the object. This is similar to increasing W for selection. Another example of using the Voronoi diagram is the Sticky Ray [Steinicke, 06] which points to the selectable object closest to the 'direction ray' of the pointer. Furthermore, the technique known as Expanding Targets [Argelaguet, 08] [Mcguffin, 02] [Mcguffin, 05] [Guillon, 15] can also improve user performance in some conditions. This technique dynamically scales the potential targets near the selection tool and provides a larger target area for the user to acquire. Although theoretically, these techniques will decrease the selection time in any condition, it is shown that they may cause confusion in densely populated spaces [Argelaguet, 08]. Other kinds of pointing facilitators based on control/displayed ratio [Frees, 07] and rank adaptation [Haan, 05] were also explored.

Prior research has considered the performance of some of the pointing facilitators in 3D virtual environments [Vanacken, 07] [Cashion, 12]. However, the level of target density has seldom been considered in detail. Furthermore, the studies were not based on HMD VR environments and could not provide the evidence that in immersive visual reality the results will be the same or similar.

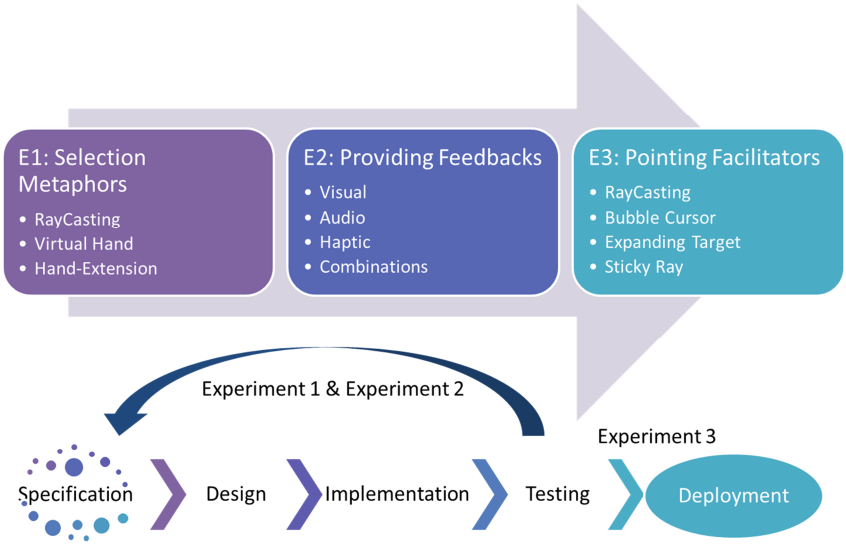

Figure 1: The software engineering process overview and the outline of our three experiments.

3 Overview of the Experiments

We propose three experiments built on each other (see Figure 1) by using the incremental development approach in software engineering. We began by designing and implementing the first experiment on selection metaphors. According to the data collected in the first experiment and user comments, we returned to the requirements and proceeded to our second experiment on providing feedback. We repeated the same flow for the third experiment on pointing facilitators. After that, we were able to distill guidelines and suggestions for the design of target selection in both sparse and dense environments.

4 Experiment 1: Selection Metaphors

With the objective of measuring the performance across different selection metaphors, in this first experiment, we compared the selection times and error rates of RayCasting and two Virtual Hand techniques (which trigger selection either by grabbing gesture or pressing the button). To distinguish the two virtual hand techniques, we named one as Virtual Hand and the other Hand-Extension (the first technique triggers selection by a *grabbing* gesture, while the second by pressing the button). The three techniques are shown in Figure 2. We were mainly concerned with finding out the best performing technique(s) among these three based on selection area and target density in 3D virtual environments in terms of selection time and error rate. To compare Virtual Hand and Hand-Extension with RayCasting, all the objects were generated close to the user (within arm's reach). Also, it is important to note that in this experiment, the defined selection time was measured on the complete procedure for the users to select the targets (including the two sub-tasks of positioning and triggering). This design would allow the apparatus themselves to maintain their selection properties (specially to distinguish the two virtual hand techniques), instead of using the same confirming technique (for example, all by pressing the button) which may not be practical for the technique in real use.

4.1 Apparatus

The study was conducted on an Intel Core i7 processor PC with an NVIDIA GTX 1080 GPU. The program was developed using C# .NET and was run on the Unity 3D platform.

Figure 3 shows the three VR devices used in this experiment. Figure 3a shows the Oculus RIFT Development Kit 2 (Oculus DK2), a lightweight headset which allows the user to step into the 3D environment and look at any direction. We drew virtual rays in this head-mounted device to assist the users during the target selection tasks (Figure 3a). The selection was triggered by the user pressing a button on the keyboard. Figure 3b displays the same Oculus headset but equipped with a Leap Motion tracker. The Leap combines infrared LEDs and two cameras to allow tracking finger

movements of users when their hands are over or in front of the sensor. The tracker was used to simulate the virtual hand metaphor. One selection was confirmed when the system detected the grab action. We used the HTC Vive headset with its controller (shown in Figure 3c) to simulate the hand-extension metaphor which triggers selection by pressing the front button. The HTC Vive controller allows the user to wirelessly interact with the virtual world. It features 24 sensors, a multi-function trackpad, and a dual-stage trigger. Grabbing can also be easily performed by this device.

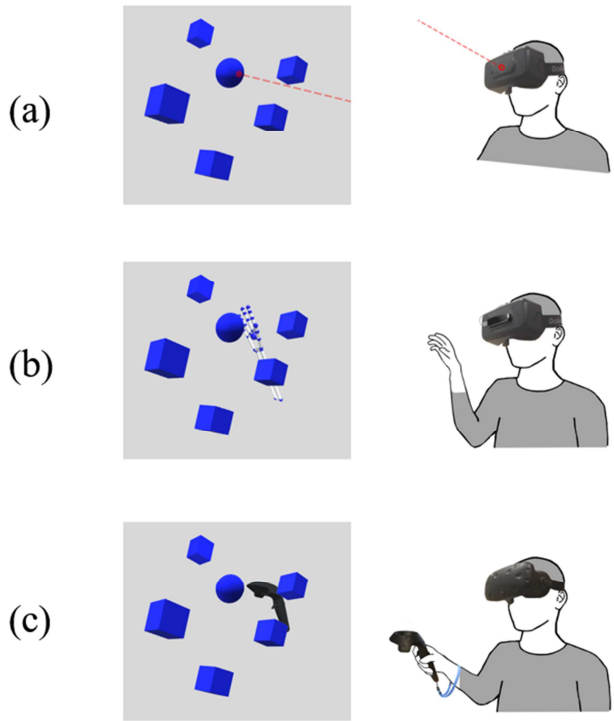

Figure 2: The technique using scenarios: (a) RayCasting; (b) Virtual Hand; and (c) Hand-Extension.

Figure 3: The three types of devices used in the first experiment: (a) the Oculus RIFT DK2; (b) the Oculus RIFT DK2 with Leap Motion tracker; and (c) the HTC Vive headset with its Dual-Handle controller.

4.2 Participants

Fifteen participants aged between 20 and 28 (Mean=23) were recruited from a local university to participate in this study. All of them are right-handed and use the right hand to control the input devices in the experiment. Seven of them have some VR experiences before the experiment.

4.3 Procedure

The whole experiment lasted approximately 45 minutes for each participant. Before the experiment started, participants were asked to complete a pre-experiment questionnaire to collect information about demographics and previous experiences with VR. After, participants were allowed to complete 9 practice trials for each device in order for them to get familiar with each technique. They were able to ask any questions during this time. Participants then completed a set of test trials. In each trial, they were presented with a set of blue cube distractor targets and one sphere goal target (see Figure 4). The positions of the goal targets were randomly generated in the scene which was divided into 9 sections from a 3×3 grid. Distractors were also randomly placed around the goal target, but based on predefined ranges of distance from it. To proceed to the next trial, participants had to successfully select the goal target. All participants were instructed to select the goal target as quickly and accurately as possible.

After finishing all the trials for one device, participants were given a 5-minute break. The experiment would end when the participant completed using all three devices—the order of the device was determined using a Latin Square approach to avoid carry-over effects. A post-experiment questionnaire was given at the end for participants to provide subjective feedback about the selection techniques.

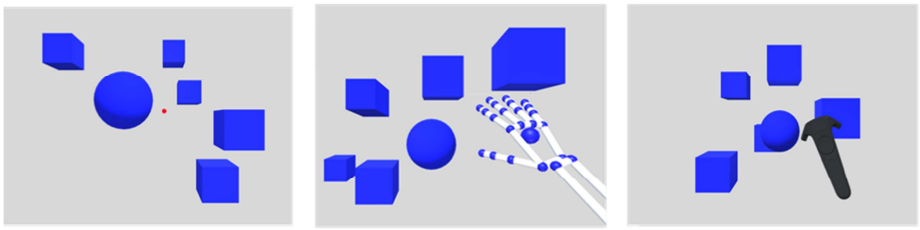

Figure 4: Screenshots of VR environment: (a) RayCasting Selection scene with 5 distractors; and (b) Virtual Hand Selection scene with 5 distractors; and (c) Hand-Extension Selection scene with 5 distractors.

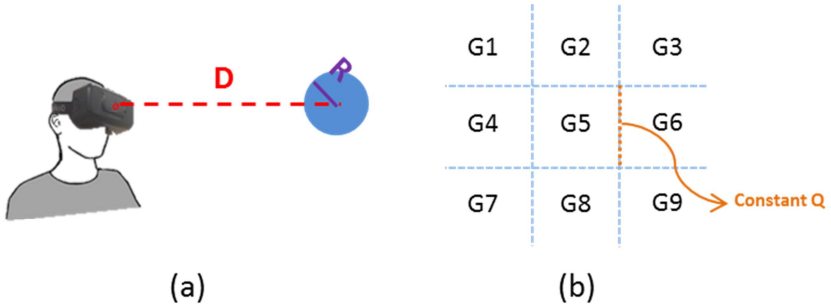

Figure 5: (a) the selection area was calculated through the distance 'D' between the central grid and the user and the radius 'R' of the goal target. (b) the position of 9 grids and each gird's length was set to a constant 'Q'.

4.4 Experiment Design

The experiment employed a 3×3×3 within-subject factorial design. The independent variables were Technique (RayCasting, Virtual Hand, and Hand-Extension), Density Level (5, 10, 15 Distractors), and the Selection Area (1, 2, 4 bits).

For Density Level (DL), we considered the number of distractors surrounding the target within a certain range of distance. In this condition, we set the maximum range of the distractor to be three times the length of the target diameter to the center of the sphere and the minimum range to be ½ of the target diameter to the center of the sphere. We also set the edge length of the distractor to be the same as the diameter of the target.

For Selection Area (SA), as shown in Figure 5a, we carefully calculated through the distance 'D' between the central grid and the user and the radius 'R' of the goal target by using the equation SA = log(D/R + 1). We set the selection area to be 1, 2, 4 and set all the objects to be within the motor spaces which could be directly reached by the users. Figure 5b shows the arrangement of the 9 grids.

Each condition would be run for 9 times in the 3×3 grid where all targets were

placed randomly in the cells. Each Participant would perform two times each combination. Thus, for 15 participants a total of 7,290 (3×3×3×9×2×15) trials were recorded. A total of 150 trials (~2%) were dropped as outliers (they had scored more than three standard deviations from the mean), leaving 7,140 logged trials.

Since we mainly focused on investigating the effectiveness of selection mechanisms, both selection times and error rates were included as dependent variables. Selection times were logged every time if the goal target was hit. An error was recorded when participants unsuccessfully attempt (e.g., the pointer was on distractor) to select the goal target (e.g., by clicking the button for the HTC Vive device). The program would detect all the errors per trial and aggregate them together.

4.5 Results

4.5.1 Selection Time

Selection time was the average time required to hits targets in a given condition. Mean times for three mechanisms are shown in Table 1.

Metaphor	Selection Time (s)	Error Rate (%)
RayCasting	2.38	7.90
Virtual Hand	2.33	7.88
Hand-Extension	2.62	13.93

Table 1: Overall means for movement time (s) and error (in %) in Experiment 1; the lower the better for both measures. RayCasting and Virtual Hand had similar performance and Hand-Extension led to longer Selection Time and higher Error Rate on average.

We analyzed the data using a 3 × 3 × 3 (Technique × DensityLevel × SelectionArea) ANOVA tests and the results showed that there was a significant effect for Technique on Selection Time ($F_{2, 28} = 62.143$, $p<.001$). Post-hoc Tukey tests revealed that both RayCasting and Virtual Hand were significantly faster than Hand-Extension ($p<.001$). There were no significant differences between RayCasting and Virtual Hand ($p>.1$).

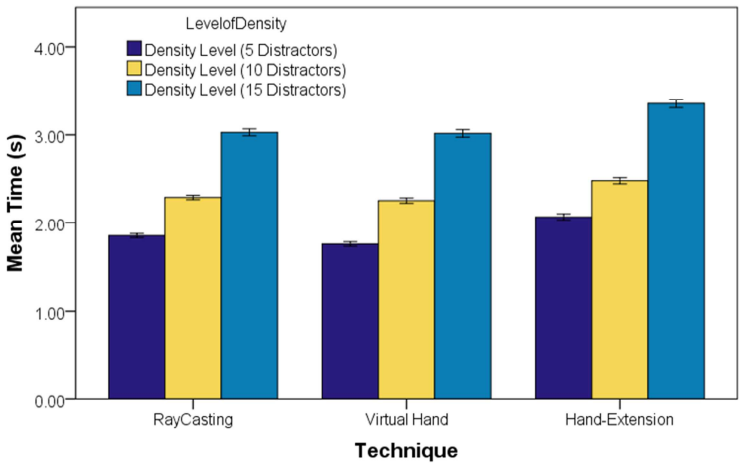

Figure 6: Mean Selection Time(s) across all three techniques. RayCasting and Virtual Hand had similar performances and were both significantly better than Hand-Extension. Increase in target density led to an increase in the time needed for all three techniques.

We also found that for Density Level there was a significant effect ($F_{2, 28}$ = 1002.446, p<.001) on Selection Time. As indicated in Figure 6, when the number of distractors went up, so did the average Selection Time. In the Low-Density condition, where the number of distractors was five, post hoc Tukey analysis found Virtual Hand was significantly faster than RayCasting (p<.05) and Hand-Extension (p<.001). In the Middle-Density condition (Number of Distractors = 10), RayCasting and Virtual Hand had no significant difference (p>0.1) on Selection Time. Both of two mechanisms were significantly faster than Hand-Extension (p<.001). The High-Density condition had 15 distractors around the target object, and as such the occlusion level was normally higher. We found no significant difference between RayCasting and Virtual Hand (p=1). To summarize this part, we found that Virtual Hand performed significantly faster in sparsely dense environments (Low-Density Level) than RayCasting and Hand-Extension. However, in denser environments, Selection Time of both RayCasting and Virtual Hand metaphor tended to be similar.

The results showed that there was no significant effect of Selection Area on Selection Time ($F_{2, 28}$ = 0.23, p>.1). However, a significant interaction effect was found between Selection Area × Number of Distractors ($F_{3, 42}$ = 8.542, p<.001).

4.5.2 Error Rate

Error trials were trials where the first click did not result in selecting the goal target. The error rates shown in Table 1 above were the mean percentages of error trials for each technique. Participants tended to make fewer errors in RayCasting and Virtual Hand than Hand Extension. ANOVA tests showed that Technique had significant effects on error rate ($F_{2, 28}$ = 32.64, p<.001) and post hoc analysis indicated that error rate for Hand-Extension was significantly higher (p<.001) than the other two

techniques. Figure 7 shows the mean error rate of Technique based on the three different numbers of distractors conditions. We found that when there were 15 distractors, the error tends to be much higher than other conditions. This might be due to the objects occluding each other, which made it very difficult for users to select the right target directly. The user also commented in the post-experiment questionnaire that when the targets got occluded with each other, they did not have an idea or clue on which targets to select if no feedback was provided.

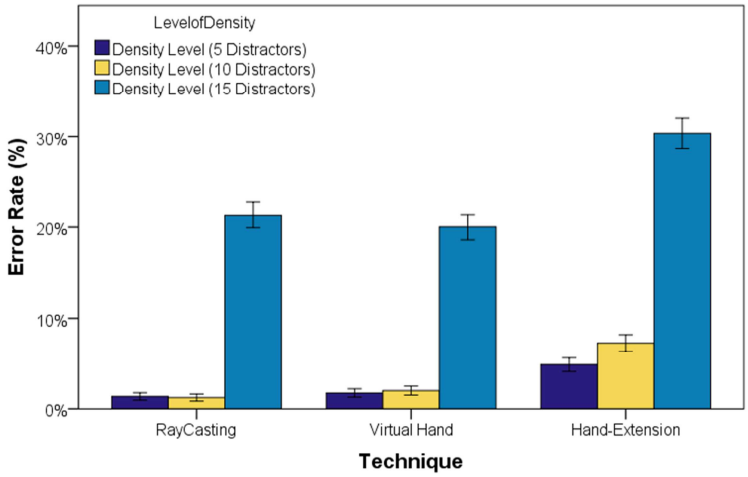

Figure 7: Error rates (%) across all three techniques. A large increase in errors is observed when there were 15 distractors.

4.6 Discussion – Experiment 1

After seeing the results, we were particularly interested in why the two virtual hand metaphors performed so differently (the virtual hand technique was significantly faster with much fewer errors) during the experiment. Since the differences between them were only the triggering technique (by grabbing or pressing the button) and their appearances in the virtual environment (the "hand" shape and the Vive Controller shape). We thought at first the slower selection for the Hand-Extension was caused by a much higher error rate during the selection tasks. We then tried to explore why this high error rate happened to Hand-Extension and we found the two following reasons.

1. Participants did not always know if they were selecting the right targets when using Hand-Extension. For example, "I thought I was on the right target but the program didn't proceed after the selection" was a participant's complaint about using Hand-Extension. Our observations are summarized in Figure 8a which shows how this inconsistency might have happened. When the goal targets were occluded by the distractors, the participants thought they might be on the right target but they were actually not. They tended to perform the selection several times in this condition, thus causing a higher error rate. However, this might not

be the case for the Virtual Hand since the grabbing gesture allows a much larger possible activation area as shown in Figure 8b.

2. *Participants tended to move from their original selection position after pressing the front button*. When observing the participants during the experiment, we noticed that by pressing the front button of the Vive Controller (the triggering process), they tended not to maintain the same position as they positioning the controller (see Figure 8c), especially for those who punched the button very hard. This might cause a higher error rate during the selection process using Hand-Extension.

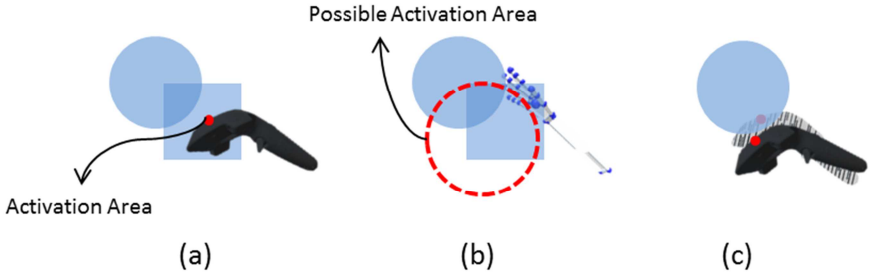

Figure 8: (a) the user performing the selection task thought the controller was on the target when using Hand-Extension; (b) virtual Hand had larger possible activation area; and (c) the selection position was originally on the target but off the target when pressing the button.

According to the above two reasons, it is not difficult to understand the differences in performance between the two techniques. This problem here might due to the different usage of the devices. However, we thought that highlighting this difference is important for designing real applications, in which the same issues would happen. These problems could be solved when the participants know which targets they are selecting. Providing additional feedback is a good choice.

5 Experiment 2: Providing Feedback

The second experiment focused on feedback for improving performance. We added visual, haptic, audio feedback and their combination to RayCasting and Virtual Hand (See Figure 9)—the best performing techniques from Experiment 1. We were particularly interested in exploring which feedback type(s) or their combinations were effective for target selection. Note that in this experiment, we still considered the selection as the whole process, including positioning and triggering, for which we measured the selection time. The rationale was because we wanted to maintain the property of the original devices.

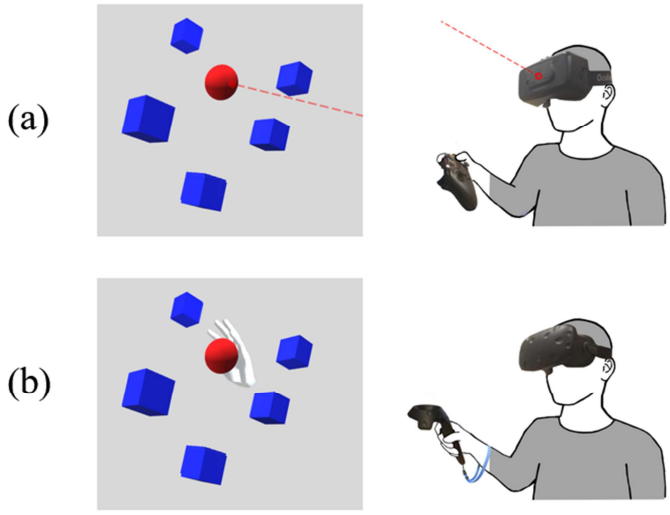

Figure 9: The technique using scenarios in Experiment 2: (a) RayCasting; (b) Virtual Hand.

5.1 Apparatus

In this experiment, we used the same PC setup as the previous experiment. For the input devices, we used the HTC Vive controller to simulate the Virtual Hand metaphor (triggering selection by pressing the button) because it was impractical to use the Leap Motion to provide vibratory feedback through the device. Users could perform grasp gestures by pressing either the Hair Trigger or the Grip Button on the Vive controller (Figure 10). The virtual hand (see Figure 9) would also perform a grasping gesture. This was more aligned to the Virtual Hand Technique instead of Hand-Extension Technique since we changed the triggering button and increased the potential activation area. For Virtual Pointing, we used the Xbox 360 controller. By pressing the X button a selection can be made.

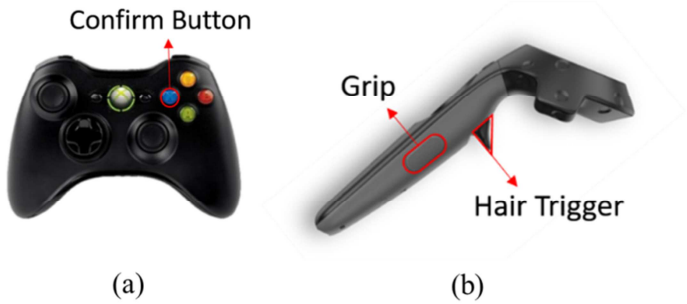

Figure 10: (a) the Xbox 360 Controller and its selection confirmation button; (b) the Grip Button and Hair Trigger of HTC Vive Controller.

5.2 Participants

Twelve participants between the ages of 22 to 28 (Mean=24) were recruited from the local university to participate in this study. All participants were right-handed. None of them had color blindness issues or have a problem with hearing. None of them participated in the first experiment.

5.3 Procedure

Before the trials started, the participants were given time to familiarize themselves with the two input mechanisms and corresponding feedback types. As with the first experiment, participants could ask any questions during this period. After this, they would proceed to do the experiment. During each trial, if the reticle of the RayCasting technique is pointing to the distractors or the goal target, or the virtual hand is touching these objects, feedback according to each condition was given to the participants. For visual feedback, the target will change to a different color. When audio feedback was triggered, a short sound would be played through the earphone attached to the VR HMD. Haptic feedback was provided through the vibrations of the HTC Vive or Xbox controller. To confirm the selection of the target, participants would press the corresponding button for either technique. The experiment ended when the participants completed both techniques with 7 feedback modes.

5.4 Experiment Design

Our experiment used the following independent variables: Technique (RayCasting and Virtual Hand) and Feedback Type (visual, audio, haptic, visual-audio, visual-haptic, audio-haptic, and visual-audio-haptic). We used a 2×7 within-subjects design in this experiment. The Selection Area was chosen to be 2 bits and number of distractors was set to be 5 in this experiment. For every condition, distractors and goal targets will be placed randomly in four quadrants or at the origin point. The above trials were completed 2 times by each participant. Therefore, for 12 participants, a total of 2×7×5×2×12 = 1,680 trials were logged. We removed 11 outliers (~0.6%) representing scores with more than three standard deviations from the mean—thus we were left with 1,669 trials. Like the previous experiment, selection time and errors were recorded.

5.5 Results

5.5.1 Selection Time

By using a 2×7 (Technique×FeedbackType) ANOVA, we found that Feedback Type had a significant effect on Selection Time ($F_{6, 66} = 18.36$, $p<.001$).

For single feedback modes, post-hoc Tukey tests showed that visual feedback was significantly faster ($p<.001$) than haptic and audio feedback. Haptic feedback and audio feedback led to the non-significant difference ($p=1$) in Selection Time. We found no significant differences when the feedback types were combined (i.e., visual-audio, visual-haptic, audio-haptic and visual-audio-haptic; ($p>.1$)).

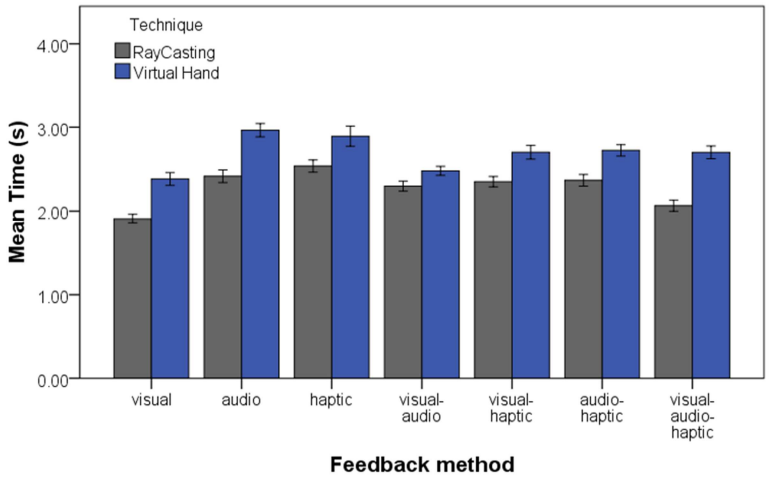

Figure 11: Mean Selection Time for RayCasting and Virtual Hand based on the 7 feedback modes. Visual feedback led to the fastest performance overall.

5.5.2 Error Rate

In this second experiment, the overall error rate for RayCasting is 0.5% and the Virtual Hand is 1.9% and a significant main effect of Technique ($F_{1, 11}$= 23.049, p<.001) was found. We also noticed that different feedback had a significant influence on the error rate ($F_{6, 66}$ = 7.47, p<.001). The error rate for visual feedback is 3.3%, audio feedback is 6.7%, haptic feedback is 10.8%, visual-audio feedback is 0.4%, visual-haptic feedback is 4.6%, audio-haptic feedback is 0.8% and visual-audio-haptic feedback is 5.0%. Post hoc Tukey showed that only using haptic feedback is significantly slower than all the other feedback types.

6 Experiment 3: Pointing Facilitators

In the third experiment, we compared the best performing selection technique with the feedback we found in the previous experiment, which was RayCasting with visual feedback only, with the current state-of-art pointing facilitators. We measured selection time and learning the effect of these techniques. Our purpose was to find the most easy-to-learn and best-performing technique for target selection in VR environments.

6.1 Apparatus

We used the same desktop settings as in the previous two experiments. Oculus Rift CV1 was used to provide a higher level of immersion. A Dell MS111-L mouse was used as the only input device for clicking and selecting targets.

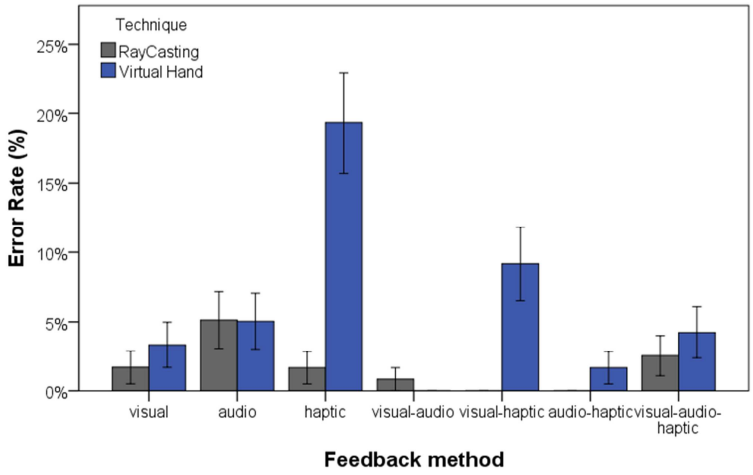

Figure 12: Mean Error Rate for RayCasting and Virtual Hand based on the 7 feedback modes.

6.2 Participants

Twelve participants between the ages of 22 to 28 (Mean=24) were recruited from the local university to participate in this study. All participants were right-handed. None of them had color blindness issues or have a problem with hearing. None of them participated in the first experiment.

6.3 Procedure

As with the previous experiment, participants completed a short questionnaire to collect demographic data and their prior experiences with VR systems. Unlike the previous two experiments, participants were not given time to get familiar with 4 selection techniques in the experiment since we want to compare the learning effect of each technique. The tests started when the participant said he or she was ready. They were asked to complete the experiment using four techniques for one block. Between each block, participants had a couple of minutes to rest before continuing to the next block. For each technique, distractors and the goal target were randomly placed in a 3×3 grid. The goal target was marked as gray and distractors were marked as blue (see Figure 13 for four scenarios). When the pointer was on a target, it would be highlighted visually, turning it into a red object. Pink was used to represent the reticle (see Figure 13). Techniques were counterbalanced according to the Latin Square design. After they completed the whole experiment, participants were asked to complete a post-questionnaire to collect their subjective preferences for the four techniques.

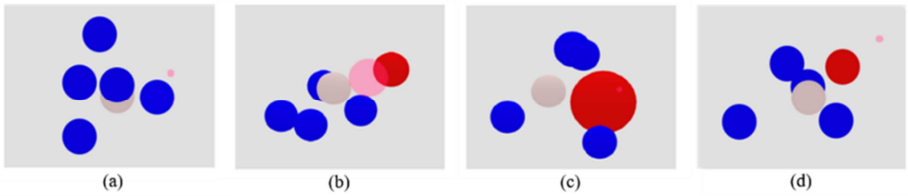

Figure 13: Four selection scenarios with 5 distractors: (a) RayCasting with visual feedback scenario; (b) Bubble Cursor scenario with one distractor under selection; (c) Expanding Target scenario with one distractor being expanded; and (d) Sticky Ray scenario with one distractor under selection.

6.4 Experiment Design

This third experiment followed a 4×4×3 within-subjects design with the following factors: Blocks (4 blocks to test learning effect), Techniques (RayCasting Selection with visual feedback, Bubble Cursor, Expanding Target, Sticky Ray), Number of Distractors (5, 10, 15). As there were 9 (3×3 grid) different locations for each condition and we had twelve participants, a total of 5,184 (4×4×3×9×12) trials were logged. As with the previous experiments, by removing 31 trials (~0.5%) as the outliers, we were left with 5,153 trials. For each randomly generated trial, the system would test if the goal target was in front or behind a distractor. The overlapping (or occluded) condition would be marked. As the previous two experiments, we recorded the selection time and errors.

In our Bubble Cursor implementation, we first sorted the distances between the cursor and all the objects in ascending order. We then chose two objects which are closest to the cursor and set the Intersecting Distance (the length of the shortest line connecting the center of the bubble cursor and the second closest object border) and Containment Distance (the length of the longest line connecting the center of the bubble cursor and the closest object border). We set the radius of the bubble cursor to a minimum of Containment Distance and Intersecting Distance [Pavlovych, 09]. When a part of the object is included in the cursor, it would be highlighted and, in such cases, we made sure that there would only be one highlighted object.

For Expanding Target technique, the object would be scaled up to make it bigger when the cursor was located within a certain range around the object (Figure 13b). We implemented Sticky Ray by selecting the closest object to the current position of the cursor.

6.5 Results

6.5.1 Selection Time

In this experiment, we used a 4×4×3 (Technique × Block × NumberOfDistractors) ANOVA. We found a significant effect of Techniques ($F_{3, 33} = 7.73$, $p<.001$), and Blocks ($F_{3, 33} = 32.431$, $p<.001$) on Selection Time. There was also a significant interaction effect of Techniques × Blocks interaction ($F_{8, 88} = 4.58$, $p<.001$). Moreover, we also found a significant effect between occluded and non-occluded

trials ($F_{1, 11} = 50.01$, $p<.001$) on Selection Time. On the other hand, the results showed did not show density to have a significant effect on Selection Time ($F_{2, 22} = 1.26$, $p>.1$).

Post hoc Tukey tests showed that only the Bubble Cursor was significantly faster ($p<.001$) than RayCasting with visual feedback. Bubble Cursor, Expanding Target, and Sticky Ray did not seem to show significant differences. However, we found that there was no significant difference ($F_{3, 33} = 1.66$, $p>.1$) among the four selection techniques on Selection Time for trials in the Overlapped condition. As shown in Figure 14, the mean Selection Time of the four techniques tends to be the same in this Overlapped condition. RayCasting was still significantly slower ($p<.001$) than Bubble Cursor in the Non-Overlapped condition.

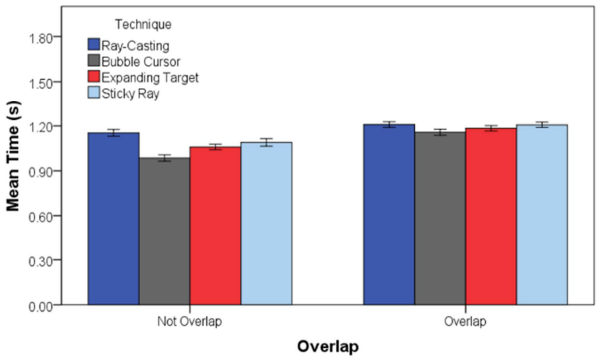

Figure 14: The mean Selection Time of the four techniques in Overlapped and Non-Overlapped conditions.

6.5.2 Error Rate

In this experiment, we did not find a significant effect of Techniques on Error Rate ($F_{2, 22} = 0.20$, $p>.1$). The RayCasting had an error rate of 5.7%, the Bubble Cursor 5.8%, the Expanding Target 6.0%, and the Sticky Ray 5.3%. We also noticed that error rate a significant difference among three density levels ($F_{2, 22} = 5.11$, $p<.05$**).**

6.5.3 Learning Effect

Learning effect seemed to be different among the four techniques. Blocks have significant effect on RayCasting with visual feedback ($F_{3, 33} = 37.34$, $p<.001$), Bubble Cursor ($F_{3, 33} = 31.63$, $p<.001$), Expanding Target ($F_{3, 33} = 7.94$, $p<.001$) and Sticky Ray ($F_{3, 33} = 7.91$, $p<.001$). For RayCasting with visual feedback, there was a significant difference ($p<.001$) between Block 1 and the other three blocks. There was no significant different ($p>.1$) among the last three blocks. As for Bubble Cursor, Block 1 was significantly faster ($p<0.01$) than Block 2 while Block 2 was significantly faster ($p<0.001$) than Block 3. There was no significant difference between Block 3 and Block 4 ($p>.1$). With Expanding Target, there was no significant difference among Blocks 1, 2, 3 ($p>.1$). These three blocks performed significant slower ($p<.05$) than Block 4. Sticky Ray had a significant increase in speed in Block 2

($p<0.01$), but no significant difference among the last three blocks ($p>.1$). Figure 16 shows the mean Selection Time of the four techniques according to blocks.

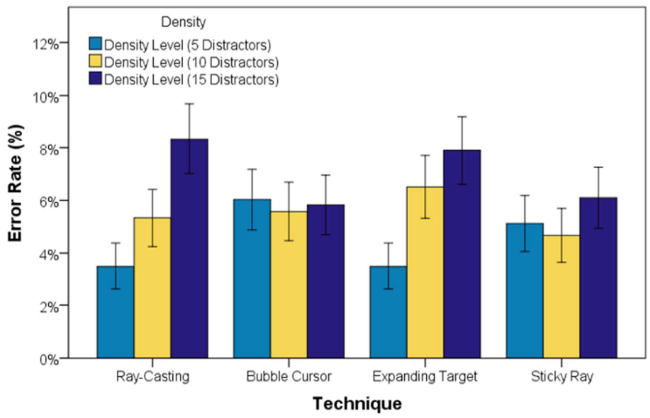

Figure 15: Error Rates (%) across all four techniques.

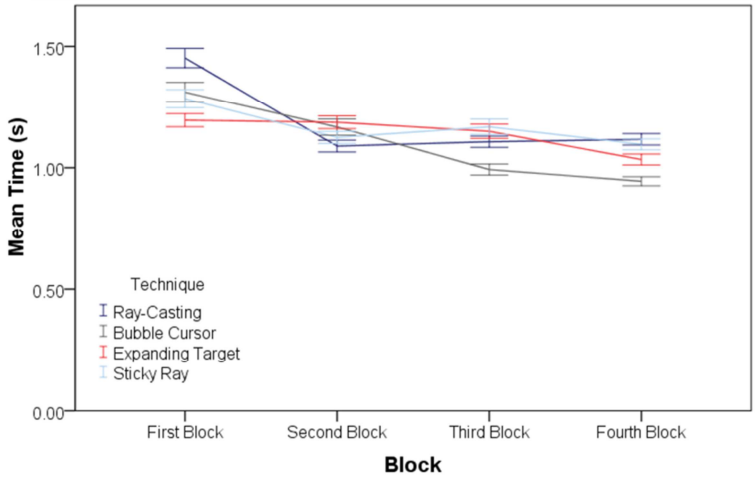

Figure 16: The mean Selection Time of the four techniques across different blocks.

6.6 User Preference

Data collected from the post-experiment questionnaire indicated that participants preferred Bubble Cursor Selection (8 of 12) than the other three techniques mostly because it was faster. On the other hand, most participants (7 of 12) indicated that RayCasting with visual feedback was the easiest to learn and more comfortable compared with the others. 8 participants commented that pointing facilitators make

interaction somewhat "*annoying*", "*confusing*" and "*unexpected*". When we asked them in what cases this was the issue, these participants said cases where the target goal was occluded and the targets were close to each other. A number of them suggested that in these two cases, it would be better to interact using non-pointing facilitators.

7 Summary and Conclusions

This research has explored target selection in HMD VR environments. We wanted to assess the performance of the main selection metaphors/techniques under conditions that were relevant to these VR environments, including a various index of difficulty (derived from the Fitt's Law), target density, and target occlusion. We also studied if enhancements to these techniques can improve performance, lower errors, and lead to better usability. To this end, we conducted three experiments.

In the first experiment, we compared the performance of three selection metaphors: RayCasting, Virtual Hand, and Hand Extension and found that RayCasting and Virtual Hand had similar performance across different conditions. Hand-Extension led to a significantly slower performance and higher errors than RayCasting and Virtual Hand. We also noticed that the three metaphors had sharp increases on error rate for cases with high-density of targets. To see if we could lower the error rates, we added various feedback modes to aid the selection task for the two best performing techniques.

In our second experiment, we added visual, audio, and haptic feedback and their combinations to the RayCasting and Virtual Hand. We found that providing only visual feedback led to improved time performance for both techniques. The results also showed that combining feedback types may not have significant improvement in the performance. Our findings seemed to be aligned with a pilot research by [Mcguffin, 02], where they found that increasing the amount/types of feedback may even reduce performance. We concluded in this experiment that providing visual feedback is likely to be sufficient.

Three pointing facilitators were used in the third experiment to assess if they could augment user performance and lower error rates. We found that Bubble Cursor was significantly better in reducing the selection time, but only in non-occluded situations. Expanding Target and Sticky Ray also improved performance significantly compared to RayCasting with visual feedback, but also only in non-occluded trials. For occluded cases, the movement time of the four techniques was comparable. In fact, participants commented that pointing facilitators may be annoying in dense conditions. This was the case especially for the Expanding Target technique, where objects became too ambiguous to capture when the cursor would move around. We found the learning effect that RayCasting technique without any pointing facilitators may be the easiest for users to learn. The results suggested that pointing facilitators which require dynamically changing the size of the cursors (e.g., Bubble Cursor) or of the targets (e.g., Expanding Target) may require a longer time for learning than other facilitators (e.g., Sticky Ray).

According to the three experimental results, we extrapolate the following lessons:

L1. Both RayCasting and Virtual Hand can provide similar performance and lead to near-equal error rates.

L2. All three metaphors (RayCasting, Virtual Hand, and Hand Extension) suffer when the density of targets increases.

L3. Visual feedback seems to be the most natural and can lead to better performance and lower error rates. The addition or combination of other feedback types does not seem to help increase performance but can lower error rates.

L4. Techniques enhanced with pointing facilitators can improve performance (when compared to RayCasting with visual feedback) but only in situations where the target is not occluded (or in high-density situations). In situations where it is occluded, techniques with pointing enhancements do not appear to have an advantage.

L5. Based on our participants' subjective feedback, it appears that pointing facilitators are preferred only when efficiency and speed of selection are required but can also lead to uncomfortable interaction.

L6. In cases where users have a very short time to familiarize themselves with the techniques, simple RayCasting with visual feedback could work best as it is aligned with how the standard cursor movement works. Pointing facilitators are not common in systems familiar to users and they will likely require some time for learning how to interact with the techniques and to adjust their mental models.

From these lessons we can further extrapolate the following design recommendations:

R1. A simple technique, like RayCasting with direct visual feedback, can work well for complex environments where many target distractors are clustered together (i.e., occlusion).

R2. When considering feedback, visual feedback seems to elicit a quick response; other types of feedback like audio and haptic can complement visual feedback but does not necessarily increase performance.

R3. A technique with enhanced pointing facilitators, like Bubble Cursor, can work well for simple environments with target distractors that are distributed sparsely.

R4. When techniques with pointing facilitators are provided, there is likely a need for users to learn first.

In short, our results support simple techniques and their uses in dense environments where occlusion often takes place—e.g., using RayCasting with visual feedback. In environments with sparse targets, pointing facilitators like Bubble Cursor can be considered to improve quick selection. We hope the contributions of this work will aid VR designers to create more usable and efficient 3D interfaces and will offer a route to developing new efficient target selection techniques for HMD VR environments.

Acknowledgments

We would like to thank the participants for their time. We would like to also thank the reviewers for the comments that helped improve our paper. This research was partially funded by the XJTLU Key Program Special Fund (KSF-P-02) and XJTLU Research Development Fund.

References

[Argelaguet, 08] Argelaguet, F., Andujar, C: Improving 3D selection in VEs through expanding targets and forced disocclusion. International Symposium on Smart Graphics. Springer, Berlin, Heidelberg, 45-57, 2008.

[Argelaguet, 09] Argelaguet, F., Andujar, C.: Efficient 3D pointing selection in cluttered virtual environments. IEEE Computer Graphics and Applications 29, 6, 34-43, 2009.

[Argelaguet, 03] Argelaguet, F., Andujar, C: A survey of 3d object selection techniques for virtual environments. Computers & Graphics 37, 3 (May 2013), 121-136, 2003.

[Arsenault, 00] Arsenault. R, Ware, C: Eye-hand co-ordination with force feedback. Proceedings of the SIGCHI conference on Human Factors in Computing Systems. ACM Press, New York, NY, 408-414, 2000.

[Akamatsu, 95] Akamatsu, M., MacKenzie, I. S., & Hasbroucq, T.: A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. Ergonomics, 38(4), 816-827, 1995.

[Bowman, 04] Bowman, D., Kruijff, E., LaViola Jr J, J., Ivan P. Poupyrev: 3D User Interfaces: Theory and Practice. Addison-Wesley, 2004.

[Bowman, 97] Bowman, D., Hodges, L: An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments. In Proceedings of the 1997 symposium on Interactive 3D graphics. ACM Press, 35-ff, 1997.

[Bowman, 01] Bowman, D., Johnson, D., Hodges, L.: Testbed evaluation of virtual environment interaction techniques. Presence: Teleoperators and Virtual Environments 10, 1 (Mar. 2001), 75-95, 2001.

[Cashion, 12] Cashion, J., Wingrave, C., LaViola Jr J, J.: Dense and dynamic 3D selection for game-based virtual environments. IEEE Transactions on Visualization and Computer Graphics 18, 4 (Apr. 2012), 634-642, 2012.

[Cockburn, 11] Cockburn, A., Quinn, P., Gutwin, C., Ramos, G., Looser, J.: Air pointing: Design and evaluation of spatial target acquisition with and without visual feedback. International Journal of Human-Computer Studies 69, 6, 401-414, 2011.

[Cournia, 03] Cournia, N., Smith, J., Duchowski, A: Gaze-vs. hand-based pointing in virtual environments. In CHI'03 extended abstracts on Human factors in computing systems. ACM Press, 772–773, 2003.

[Haan, 05] Haan, G., Koutek, M. and Post, F.: IntenSelect: Using Dynamic Object Rating for Assisting 3D Object Selection. In IPT/EGVE, 201-209, 2005.

[Ebrahimi, 16] Ebrahimi, E., Babu, S., Pagano, C., and Jörg, S.: An empirical evaluation of visuo-haptic feedback on physical reaching behaviors during 3D interaction in real and

immersive virtual environments. ACM Transactions on Applied Perception (TAP) 13, no. 4 (2016): 19.

[Farmani, 17] Farmani, Y., and Teather, R.: Player performance with different input devices in virtual reality first-person shooter games. Proceedings of the 5th Symposium on Spatial User Interaction, pp. 165-165. ACM, 2017.

[Fitts, 54] Fitts, P: The information capacity of the human motor system in controlling the amplitude of movement. Journal of experimental psychology 47, 6 (Jun. 1954), 381, 1954.

[Fitzmaurice, 93] Fitzmaurice, G. W.: Situated information spaces and spatially aware palmtop computers. Communications of the ACM, 36(7), 39-49, 1993.

[Forsberg, 96] Forsberg, A., Herndon, K., Zeleznik. R.: Aperture based selection for immersive virtual environments. In Proceedings of the 9th annual ACM symposium on User interface software and technology. ACM Press, New York, NY, 95–96, 1996.

[Frees, 07] Frees, S., Kessler, G., Kay, E.: PRISM interaction for enhancing control in immersive virtual environments. ACM Transactions on Computer-Human Interaction (TOCHI) 14, 1, 2, 2007.

[Grossman, 05] Grossman, T., Balakrishnan, R.: The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. Proceedings of the SIGCHI conference on Human factors in computing systems. ACM Press, New York, NY, 281-290, 2005.

[Grossman, 06] Grossman, T., Balakrishnan, R.: The design and evaluation of selection techniques for 3D volumetric displays. Proceedings of the 19th annual ACM symposium on User interface software and technology. ACM Press, New York, NY, 3-12, 2006.

[Guillon, 15] Guillon, M., Leitner, F., and Nigay, L.: Investigating Visual Feedforward for Target Expansion Techniques. Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, pp. 2777-2786. ACM, 2015.

[Guillon, 16] Guillon, M., Leitner, F., and Nigay ,L.: Target Expansion Lens: It is Not the More Visual Feedback the Better!. Proceedings of the International Working Conference on Advanced Visual Interfaces, pp. 52-59. ACM, 2016.

[Lee, 16] Lee, J., Park, J., Oh, J., and Lee. J.: Fast and Accurate 3D Selection using Proxy with Spatial Relationship for Immersive Virtual Environments. Proceedings of the 2016 Symposium on Spatial User Interaction, pp. 209-209. ACM, 2016.

[Liang, 18] Liang, H.-N., Lu, F., Shi, Y., Nanjappan, V. and Papangelis, K.: Evaluating the effects of collaboration and competition in navigation tasks and spatial knowledge acquisition within virtual reality environments. Future Generation Computer Systems. 2018.

[Liang, 16] Liang, Hai-Ning, Yuwei Shi, Feiyu Lu, Jizhou Yang, and Konstantinos Papangelis.: VRMController: an input device for navigation activities in virtual reality environments. In Proceedings of the 15th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its Applications in Industry-Volume 1, pp. 455-460. ACM, 2016.

[Herndon, 94] Herndon, K., Dam, A., Gleicher, M.: The challenges of 3D interaction: a CHI'94 workshop. ACM SIGCHI Bulletin 26, 4, 36-43. ACM Press, 1994.

[Kulik, 09] Kulik, A.: Building on realism and magic for designing 3d interaction techniques. IEEE Computer Graphics and Applications 29, 6 (Nov.-Dec. 2009), 22–33, 2009.

[Kytö, 18] Kytö, M., Ens, B., Piumsomboon, T., Lee, G., and Billinghurst. M.: Pinpointing: Precise Head-and Eye-Based Target Selection for Augmented Reality. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, p. 81. ACM, 2018.

[Liang, 94] Liang, J., Green, M.: Jdcad: A highly interactive 3d modeling system. Computers & graphics 18, 4 (Jul.-Aug. 1994), 499–506, 1994.

[Lubos, 14] Lubos, P., Bruder, G., Steinicke, F., Analysis of direct selection in head-mounted display environments. In 3D User Interfaces (3DUI), 2014 IEEE Symposium on IEEE, 11-18, 2014.

[MacKenzie, 92] MacKenzie, I.: Fitts' law as a research and design tool in human-computer interaction. Human-computer interaction 7, 1 (Mar. 1992), 91-139.

[Mayer, 18] Mayer, S., Schwind, V., Schweigert, R., and Henze. N.: The Effect of Offset Correction and Cursor on Mid-Air Pointing in Real and Virtual Environments. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, p. 653. ACM, 2018.

[Mcguffin, 02] Mcguffin, M., Balakrishnan, R.: Acquisition of expanding targets. In Proceedings of the SIGCHI conference on Human factors in computing systems ACM, 57-64, 2002.

[Mcguffin, 05] Mcguffin, M., Balakrishnan, R.: Fitts' law and expanding targets: Experimental studies and designs for user interfaces. ACM Transactions on Computer-Human Interaction (TOCHI) 12, 4, 388-422, 2005.

[Mine, 95] Mine, M.: Virtual environment interaction techniques. UNC Chapel Hill computer science technical report TR95-018, 1995.

[Mould, 04] Mould, D., Gutwin, C.: The effects of feedback on targeting with multiple moving targets. In Proceedings of Graphics Interface 2004 Canadian Human-Computer Communications Society, 25-32, 2004.

[Papangelis, 17] Papangelis, K., Metzger, M., Sheng, Y., Liang, H.-N, Chamberlain A., and Cao, T.: Conquering the City: Understanding perceptions of Mobility and Human Territoriality in Location-based Mobile Games. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 1, no. 3 (2017): 90.

[Papangelis, 17] Papangelis, K., Metzger, M., Sheng, Y., Liang, H.-N, Chamberlain, A., and Khan V.: "Get off my lawn!": Starting to understand territoriality in location based mobile games. (2017): 1955-1961.

[Papangelis, 17] Papangelis, K., Sheng, Y., Liang, H.-N, Chamberlain, A., Khan, V., and Cao. T.: Unfolding the interplay of self-identity and expressions of territoriality in location-based social networks. Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers, pp. 177-180. ACM, 2017.

[Pavlovych, 09] Pavlovych, A., Stuerzlinger, W.: The tradeoff between spatial jitter and latency in pointing tasks. Proceedings of the 1st ACM SIGCHI symposium on Engineering interactive computing systems. ACM Press, New York, NY, 187-196, 2009.

[Pfeiffer, 15] Pfeiffer, M., Stuerzlinger, W.: 3D virtual hand pointing with EMS and vibration feedback. In 3D User Interfaces (3DUI), 2015 IEEE Symposium on IEEE, 117-120, 2015.

[Poupyrev, 96] Poupyrev, I., Billinghurst, M., Weghorst, S., Ichikawa, T.: The gogo interaction technique: non-linear mapping for direct manipulation in vr. In Proceedings of the 9th annual ACM symposium on User interface software and technology. ACM Press, 79–80, 1996.

[Poupyrev, 99] Poupyrev, I., Ichikawa, T.: Manipulating Objects in Virtual Worlds: Categorization and Empirical Evaluation of Interaction. Journal of Visual Languages & Computing 10, 1 (Feb. 1999), 19–35, 1999.

[Poupyrev, 98] Poupyrev, I., Ichikawa, T., Weghorst, S., Billinghust. M.: Egocentric object manipulation in virtual environments: empirical evaluation of interaction techniques. Computer graphics forum 17, 3, 41-52. Blackwell Publishers Ltd, 1998.

[Qian, 17] Qian, Y., and Teather, R.: The eyes don't have it: an empirical comparison of head-based and eye-based selection in virtual reality. Proceedings of the 5th Symposium on Spatial User Interaction, pp. 91-98. ACM, 2017.

[Steed, 04] Steed, A., Parker, C.: 3D selection strategies for head tracked and non-head tracked operation of spatially immersive displays. 8th International Immersive Projection Technology Workshop, 13-14, 2004.

[Stoakley, 95] Stoakley, R., Conway, M. J., & Pausch, R.: Virtual reality on a WIM: interactive worlds in miniature. In Proceedings of the SIGCHI conference on Human factors in computing systems (pp. 265-272). ACM Press/Addison-Wesley Publishing Co, 1995.

[Steinicke, 06] Steinicke, F., Ropinski, T., Hinrichs, K.: Object selection in virtual environments using an improved virtual pointer metaphor. Computer Vision and Graphics 32, 320-326, 2006.

[Teather, 11] Teather, R., Stuerzlinger, W.: Pointing at 3d targets in a stereo head-tracked virtual environment. In 3D User Interfaces (3DUI), 2011 IEEE Symposium on. IEEE, Singapore, Singapore, 87–94, 2011.

[Teather, 14] Teather, R., Stuerzlinger, W.: Visual aids in 3D point selection experiments. In Proceedings of the 2nd ACM symposium on Spatial user interaction ACM, 127-136, 2014.

[Tran, 17] Tran, T., Shin, H., Stuerzlinger, W., and Han, J.: Effects of virtual arm representations on interaction in virtual environments. Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology, p. 40. ACM, 2017.

[Vanacken, 07] Vanacken, L., Grossman, T., Coninx, K.: Exploring the effects of environment density and target visibility on object selection in 3d virtual environments. In 3D User Interfaces, 3DUI'07. IEEE Symposium on. IEEE, Charlotte, NC, 2007

[Vanacken, 06] Vanacken, L., Raymaekers, C., Coninx, K.: Evaluating the influence of multimodal feedback on egocentric selection metaphors in virtual environments. International Workshop on Haptic and Audio Interaction Design. Springer, Berlin, Heidelberg, 2006.

[Wanger, 92] Wanger, L.: The effect of shadow quality on the perception of spatial relationships in computer generated imagery. Proceedings of the 1992 symposium on Interactive 3D graphics. ACM Press, New York, NY, 39-42, 1992.

[Wingrave, 05] Wingrave, C., A. Bowman, D. A.: Baseline factors for raycasting selection. Proceedings of HCI International, 2005.

[Wingrave, 02] Wingrave, C. A., Bowman, D. A., Ramakrishnan, N., Towards preferences in virtual environment interfaces. In EGVE. 63–72, 2002.

[Zaranek, 14] Zaranek, A., Ramoul, B., Yu, H., Yao, Y., and Teather. R.: Performance of modern gaming input devices in first-person shooter target acquisition. CHI'14 Extended Abstracts on Human Factors in Computing Systems, pp. 1495-1500. ACM, 2014.