Exploring the Effects of Intended Use on Targeting in Virtual Reality

Category: Research

Abstract

Researchers have shown that distance and size are not the only factors that impact the target acquisition time in desktop interfaces, but that its intended use, whether it is selected, dragged, or otherwise manipulated, can also have a significant influence. However, despite the increasing popularity of virtual 3D environments, the intended use of targets in these contexts has never been investigated, in spite of the richer, multidimensional manipulations afforded by these environments. To better understand the effects of intended use on target acquisition in virtual environments, we present the results of a study examining five different manipulation tasks: targeting, dual-targeting, throwing, docking and reorienting. Our results demonstrate that the intended use of a target affects its acquisition time and, correspondingly, the movement towards the target. As these environments become more commonplace settings for work and play, our work provides valuable information on throughput, applicable to a wide range of tasks.

Index Terms: Human-centered computing—Virtual reality; Human-centered computing—Pointing

1 Introduction

Target acquisition is one of the most common actions performed in an interactive system regardless of whether the computing platform is a desktop computer [25, 36], a touch-based device [23], a large physical display [51] or an Augmented Reality environment [13].

Fitts's Law [36] is the most commonly used model to describe the movement time taken to acquire a target. Movement time (*MT*) for a one-dimensional pointing task is described by a linear function of the index of difficulty (*ID*) of the pointing task:

$$MT = a + b \cdot ID \tag{1}$$

where a, b are empirically determined regression coefficients and ID is a logarithmic term of target amplitude (A) and target width (W) [35]:

$$ID = log_2\left(\frac{A}{W} + 1\right) \tag{2}$$

However, target amplitude and width are not the only factors impacting acquisition performance [62]. One factor that may impact target acquisition time is the *intended use* of the target. By *intended use*, we refer to manipulations that a user intends to do or will do with the target they acquire. For example, in virtual environments, beyond simply acquiring an object, users can scale objects and move them around in the space [61], rotate objects to reveal the occluded view [15] or manipulate objects' motion to simulate physical phenomena [4]. In order to perform any of these intended uses of a target, we must first acquire it. Acting on an object in an interface is thus a compound action comprised of acquiring it and performing the action or manipulation of it (See Figure 1) and the question we pose in this paper is whether the manipulation that we will perform on an object in a virtual environment, our intended use, impacts the performance of acquiring the target.

In 2D computer interfaces, research has shown that target acquisition time may vary depending on whether the user wishes to simply acquire a target (e.g. targeting), move it (e.g. dragging or docking), or throw it (e.g. flicking) – independent of the time taken by the subsequent task [37, 47]. It is always tempting to assume that results in 2D can be directly applied to 3D, particularly to Virtual Reality (VR), but we hesitate to make such assumption for several reasons. For instance, the controller in VR is an absolute input while the mouse is a relative input, leading to different acquisition and manipulation behaviours, and some tasks more frequent in 3D were not investigated in 2D (e.g. reorienting a target). Besides, to the best of our knowledge no similar analysis has been performed in the increasingly common context of direct target acquisition in VR. As we incorporate VR into both work [17, 34] and entertainment [50], understanding the impact of various independent variables on performance measures for VR-based selection tasks remains an area of active research interest (see [9] for a review).

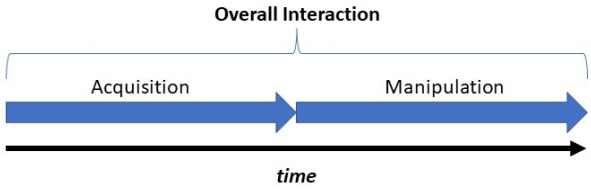

Figure 1: The overall interaction with an object includes two sequential steps: acquiring the object – time characterized by Fitts' Law or a variant – and manipulating the object – time varies depending on the complexity and nature of the manipulation.

To explore the impact of intended use on targeting in virtual environments, more precisely in the context of proximal interaction using the virtual hand metaphor, we examine five common manipulations in various VR systems, such as games [3], social [38] and educational applications [24] – Targeting (selecting a target), DualTargeting (selecting one then another target), Throwing (acquiring and pushing a target), Docking (acquiring and placing a target) and Reorienting (rotating a target about its axes) – and measure how long it takes to acquire the target to be manipulated. We find that simple targeting exhibits the fastest acquisition times for that target, reorienting the slowest acquisition times, and other intended target manipulations result in acquisition times between these extremes. Given the initial finding that the intended use of the target impacts time, we probe additional characteristics of movement toward a target with the goal of understanding how and why prior movement time is impacted by the subsequent intended use of the target. Movement profiles highlight both characteristics and differences in peak speed, time to peak speed, movement prior to selection and selection speed across different intended uses.

This paper is organized as follows. After presenting the related work on target acquisition modelling in both 2D and 3D, we describe and explain the experiment setup and task design of our study. Then, we detail our analysis from three aspects: selection time, motion kinematics and Fitts' law modelling. Finally, we present our findings and discuss their implications.

In summary, this is the first work to explore the impact of intended use on targeting in VR and our contributions include:

- Investigating impacts of intended use on target acquisition in VR and summarizing characteristics of each use.
- Presenting design implications for interfaces and interactions in VR.

2 Related Work

Fitts' Law [21, 36] is probably the most well-examined relationship in human-computer interaction (HCI) research. While Fitts' Law was originally formulated in terms of a 1D pointing task, researchers in HCI have long recognized that understanding the cost of target acquisition in graphical interfaces is useful, as it allows us to characterize the relative efficiency of different arrangements of interfaces. As a result, researchers have proposed a number of extensions to Fitts' Law to model 2D targeting [2], 3D targeting [53], gaze-based targeting [57], foot-based movement [27, 56], among others. Researchers have also explored modelling error in Fitts' Law [59], generalized Fitts' Law to incorporate steering tasks [1, 41], and leveraged the fundamental components of Fitts' Law to design a host of pointing facilitation techniques (see [6] for a review).

The goal of research into Fitts' Law is to understand and improve throughput [62] in the use of interactive systems. If we characterize the temporal cost and error rate of individual interactions (e.g. selecting a target, manipulating that target, keyboarding, homing, etc.), then the overall temporal cost and error rate of a compound interaction is the sum of these basic interactions [12, 18]. By improving the speed of target acquisition via new target acquisition techniques [6] and the speed of interactions via new interaction techniques [9], each of these subtasks becomes more efficient, increasing the overall efficiency of the interaction. Essentially, the assumption is that each individual interaction can be independently optimized. This is equally true in VR research: Bergström et al. [9], analyzing 20 years of VR research, note that studies with selection tasks measure the time a participant needs to select the next target (occasionally with additional measures of error and throughput for the selection task), while studies with manipulation tasks measure the time a participant manipulates virtual objects.

While, to the best of our knowledge, the independence of basic interactions has never been evaluated in VR, we have reason to believe that this assumption is questionable in traditional two-dimensional computer interfaces. Mandryk and Lough [37] examined how the *intended use* of a target impacts the time it takes to acquire the target. Mandryk and Lough note that, in real-world interfaces, the user acquires a target with a specific goal in mind. Perhaps they wish to click the target (targeting). Perhaps the target activates a secondary set of widgets and they need to then click on a second target (i.e. dual targeting). Perhaps they wish to move the target in some way, e.g. to re-position it imprecisely (flicking or dragging) or to re-position it precisely at a new location (docking). Mandryk and Lough found that, if the intended use was targeting or dual targeting, participants acquired the target significantly faster than if the intended use was flicking or dragging. They also noted a difference in acceleration and deceleration during the selection of the target to be manipulated: if the participants intended to flick or dock the target, then the selection movement toward that target exhibited a higher peak speed and a longer deceleration phase. Follow-on work by Ruiz and Lank [47] replicated these results and, via a more complete analysis of movement profiles, analyzed their potential impact on kinematics-linked endpoint prediction [33]. While Ruiz and Lank note that the impact on kinematics-linked modelling was likely not a concern, in both cases results implied that the potential benefits of new interaction techniques observed during manipulation may not be realized if they result in a corresponding increase in acquisition time for the target to be manipulated.

There is also a significant possibility that virtual reality manipulations will differ from both real-world and two-dimensional interface manipulations. Considering real-world manipulations, the field of psychology has actively studied the act of reaching and grasping for many years [46]. Factors, including perception of the object to be grasped [22], the manipulations to be performed on the object [49], and tactile feedback during the act of grasping [29] impact both trajectories toward an object and the positioning and speed associated with the grasping of an object. However, there exists a disconnect between real world affordances and perceived affordances [42] and an absence of the physiological interactions between hand and object, which may impact behaviour. Furthermore, while it is always tempting to assume that previous results in 2D [37, 47] can be directly applied to 3D, especially to VR, past research indicates that this may not be true. As one example of this, Cockburn and Mckenzie [16] found that user performance deteriorated for a locate-and-point task when transforming from a 2D interface to a 3D interface. Furthermore, movement planning, whether in real world or in computer interfaces, requires trajectory planning [33, 46]; in immersive VR environments, depth has been found to greatly impact both perceived width and distance [5, 14], a factor that may significantly alter the trajectory, kinematics and the impact of intended use.

3 Methodology

We conducted a controlled experiment to investigate the effect of intended use of a target on the time taken to acquire it in VR. We focused on interaction at arms' length using a 1-to-1 mapping of physical controller movement to virtual controller movement, a direct 3D target acquisition technique common in VR [14, 44, 45]. Figure 2 depicts the timeline of a trial in this experiment.

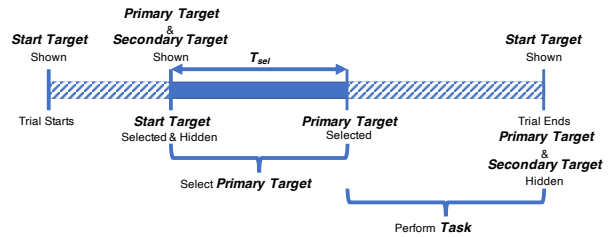

Figure 2: A timeline of a trial: A trial starts by showing the *Start Target*. Once it is successfully acquired, it vanishes and a *Primary Target* is shown, together with a *Secondary Target* if required by the *Task*. Participants acquire the *Primary Target* and perform the *Task*. *Start Target* is shown again when the trail ends.

3.1 Apparatus and Participants

The system was implemented in Unity 2020.3.7f1 with the Oculus Integration v29.0 and the study was conducted standalone on the Oculus Quest 2 (tracking frequency is 72 Hz) with Oculus left- and right-hand controllers.

In pilot studies we noted that the virtual controllers could occlude targets and their irregular shape caused problems of precision when selecting a target (because the exact selection location was unclear). To address this, we used a smaller controller 3D model (from GEAR VR), made the models translucent, and added a blue sphere (**cursor**) to the top side of the virtual controller indicating selection location, as shown in Fig. 3. Although the physical and virtual controllers are not in the same shape, this design rarely affects how participants recognize them as participants only see virtual controllers rather than physical ones. In addition, participants focus on using **cursor**, rather than the model for interactions. Making the model semi-transparent also does not impact the selection performance of a virtual hand [31, 55]. Besides, **cursor** is useful for target selection in VR [7] and our design presents better visual tracking on it.

We also noted that, as participants became fatigued, they would sometimes switch hands for a brief period, interacting with their non-dominant hand which impacted performance. Therefore, alongside instructing participants to perform the study only with their dominant hand, we disabled the non-dominant hand controller at the outset of the study and invited participants to take breaks during the experiment if needed.

A total of 15 participants, aged from 22 to 33 ($M = 26.1, SD = 3.1$, 5 identified as women, all right-handed) participated in the study. 12 participants had experienced VR prior to the study; those participants primarily used VR for entertainment activities, such as playing games. The experiment took approximately 45 minutes and participants received $15 for their participation.

3.2 Interaction for Target Acquisition

Fig. 3 illustrates the visual interaction of the target acquisition. To acquire a target, participants place the cursor (1 cm width) partially inside a target and press the selection button (in our implementation, A on the right controller) to select that target. When any part of the cursor enters a target, the target is highlighted in translucent blue, indicating that it is now selectable. A successful selection turns the cursor green. If instead the button is pressed while the cursor is outside the target, the cursor turns red and an alert sound is played to inform participants of the erroneous selection. Participants are expected to correctly acquire each target, and they can proceed after a correct selection.

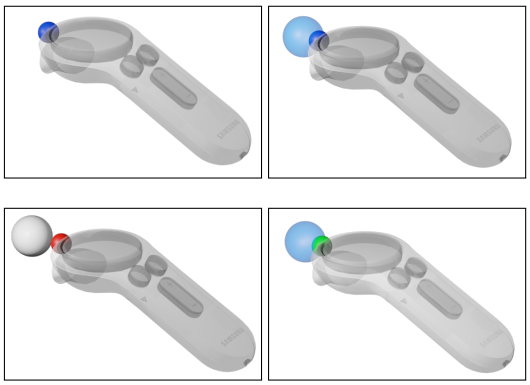

Figure 3: Input controller: when *Cursor* enters a target, the target turns translucent blue. Incorrect selections turn *Cursor* into red while correct selections turn it into green.

Given that the overall goal of this experiment was to measure how the intended use of a target impacts the time taken to acquire that target, the overall interaction requires acquisition of a target followed by a manipulation of that target, i.e. an intended use or *Task*. The interaction to start a *Task* proceeds as follows:

1. A participant acquires an initial target, the *Start Target*. To ensure that all targets in different tasks are noticeable and reachable within the field of view, the *Start Target* is positioned 35 cm in front of and 20 cm below the head.

2. When the participant correctly selects the *Start Target*, the *Start Target* vanishes and a *Primary Target* is displayed immediately. He/She is asked to acquire and manipulate this *Primary Target*. Considering that the arms of healthy adults typically reach at least 60 cm from the torso [5, 14, 30, 32, 43], targets are placed within 60 cm of participants in a region surrounding the location of the *Start Target*.

3. Participants move the cursor from the *Start Target* to the *Primary Target* and perform a *Task* on the *Primary Target*.

By default, the *Start Target*'s position remains the same in the virtual world for all tasks and targets are anchored relatively to the *Start Target*'s position, which do not follow participants' movement. However, as participants could adjust their position during the experiment, we allow them to re-calibrate the *Start Target*'s position when they notice a shift in their position, before starting any new *Task*. In practice, the re-calibration was seldom used.

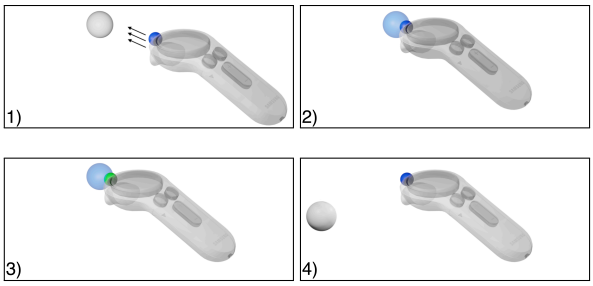

Figure 4: Correctly acquiring *Start Target* to calibrate starting position and reveal the task (the example shows TARGETING).

3.3 Independent Variables

Independent variables (IVs) in this study include the Index of Difficulty (*ID*) of the *Primary Target*, *Primary Direction*, *Secondary Direction* and *Task*.

ID of the acquisition action is computed using [35] for the Amplitude **A** (distance between the *Start Target* and the *Primary Target*) of values 9 cm, 12.5 cm, or 16 cm, and the Width **W** of the *Primary Target* of values 3 cm or 5 cm). This yields six different IDs (1.49, 1.81, 2.00, 2.07, 2.37, and 2.66 bits).

Primary Direction **D** and *Secondary Direction* $\mathbf{D}'$ (Up, Down, Forward, Backward, Left or Right) represent the direction of movement from the *Start Target* to the *Primary Target* and the *Primary Target* to the *Secondary Target* respectively.

The rationale for these values is as follows. The *Primary Target* can be located in six basic directions from the *Start Target*, i.e. Up, Down, Forward, Backward, Left and Right. Given the position of the *Start Target*, six distinct directions, and the possible existence of sequential manipulation after acquiring the *Primary Target*, the maximum values of target amplitude and target width are set to 20 cm and 10 cm respectively to avoid a target appearing outside of the field of view or outside of the reachable workspace. These constraints motivate the above values for independent variables. With this configuration, participants observe targets that are close to them from a top view rather than a straight horizontal view, which helps to reduce the depth impact on the perceived width and addresses the occlusion issue between *Primary Target* and *Secondary Target*.

3.3.1 Task

The *Tasks* that participants were asked to complete were one of five manipulations of the primary target: TARGETING, DUALTARGETING, DOCKING, THROWING, or REORIENTING (illustrated in Fig. 5). Detailed description of individual tasks follows.

TARGETING: Correctly selecting the *Start Target* of width W reveals a white sphere (*Primary Target*) of width W located at the amplitude A in direction D from the position of *Start Target*. Participants simply have to acquire the *Primary Target* by moving the controller to the target and selecting it.

DUALTARGETING: Correctly selecting the *Start Target* of width W reveals a white sphere (the *Primary Target*) and a red sphere (the *Secondary Target*) both of width W. The *Primary Target* is located in direction D with amplitude A from *Start Target* while *Secondary Target* is located in another direction D' (different from previous and next trials) with amplitude 9 cm from the *Primary Target* (a distance that guarantees that both targets still remain within arms' reach for the participant, see Sect. 3.3). Participants must select these two spheres sequentially: first the *Primary Target*, and then the *Secondary Target*.

DOCKING: Correctly selecting the *Start Target* of width W reveals a white sphere (the *Primary Target*) of width W and amplitude A in direction D, and a semi-transparent sphere (the *Secondary Target*) 1.5 times larger than the *Primary Target* and located 9 cm away

from the *Primary Target* in direction D'. The width of the *Secondary Target* reduces the required precision of the task, allowing participants to finish this task more easily. Participants are instructed to drag-and-release the *Primary Target* into *Secondary Target*.

THROWING: Correctly selecting the *Start Target* of width W reveals a white sphere (the *Primary Target*) of width W located at a position direction D from *Start Target* and at a distance of amplitude A, and a semi-transparent green wall (the *Secondary Target*) located 9 cm away from the *Primary Target* in direction D'. To reduce task difficulty (wall size being too small) and avoid visual distraction (wall size being too big), the size of this wall is set to $40\,\text{cm} \times 40\,\text{cm} \times 1\,\text{cm}$ which is over 8 times larger than the size of *Primary Target* in width. Participants are instructed to *"throw"* the *Primary Target* towards the *Secondary Target* by releasing the pressed button. The released *Primary Target* then moves in the throwing direction.

REORIENTING: Correctly selecting the *Start Target* of width W reveals a white object (the *Primary Target*) and a red object (the *Secondary Target*) whose bodies are both spheres of width W. Different shapes (i.e. a capsule, a cylinder and a cube) are placed into both *Primary Target* and *Secondary Target* to indicate the orientation, which only serve as visual references and are not selectable. The *Primary Target* is located at amplitude A from *Start Target* in direction D and *Secondary Target* 9 cm away from *Primary Target* in direction D'. The *Secondary Target* has an orientation along its roll axis with a random angle in the range ($-\frac{\pi}{2}$ to $\frac{\pi}{2}$), values empirically obtained from pilot studies to reduce clutching and task difficulty. To move to the next trial, participants were asked to rotate the *Primary Target* such that the *Primary Target* has a similar orientation to the *Secondary Target*, that is when the angle difference in each axis is smaller than $\frac{\pi}{12}$. The colour of *Secondary Target* changes to green to inform participants of the correct orientation of the *Primary Target*.

Except TARGETING, all other tasks require a manipulation sub-task on either the *Primary Target* or the *Secondary Target*. For DUALTARGETING, a second selection is required on the *Secondary Target* and other tasks require a "press and hold" behaviour on the *Primary Target* for object manipulation. While behaviours in the manipulation sub-task vary across *Task*, the acquisition time of the *Primary Target* (delimited by the controller "button down" action on the primary target) is consistent across all tasks.

3.4 Procedure

Participants were recruited from our local university. To preserve social distancing requirements, the study was conducted remotely, and both the VR headset and controllers were sanitized before being delivered to participants.

Before written consent was obtained, participants were asked to read an information letter in which they were warned about potential motion sickness as a result of wearing the head-mounted display device and that they could stop the experiment at any time without penalty if they felt uncomfortable or simply did not wish to continue. They were then asked to watch an instructional video, answer a demographics questionnaire (asking for their gender [52], age, handedness, and VR experience). They were then asked to sign the informed consent.

Participants started the experimental software by connecting with one of the researchers via video conferencing tools. The researcher verified informed consent, walked participants through disabling the non-dominant hand controller, and then guided the participant to start the experiment. In addition to researchers' verbal guidance, participants could also follow the instructions text in the system during the experiment, which presented detailed steps to guide them to control the system and finish each task.

As noted above, to evaluate the impact of intended use on target acquisition in VR, we represented different intended uses as a *Task*

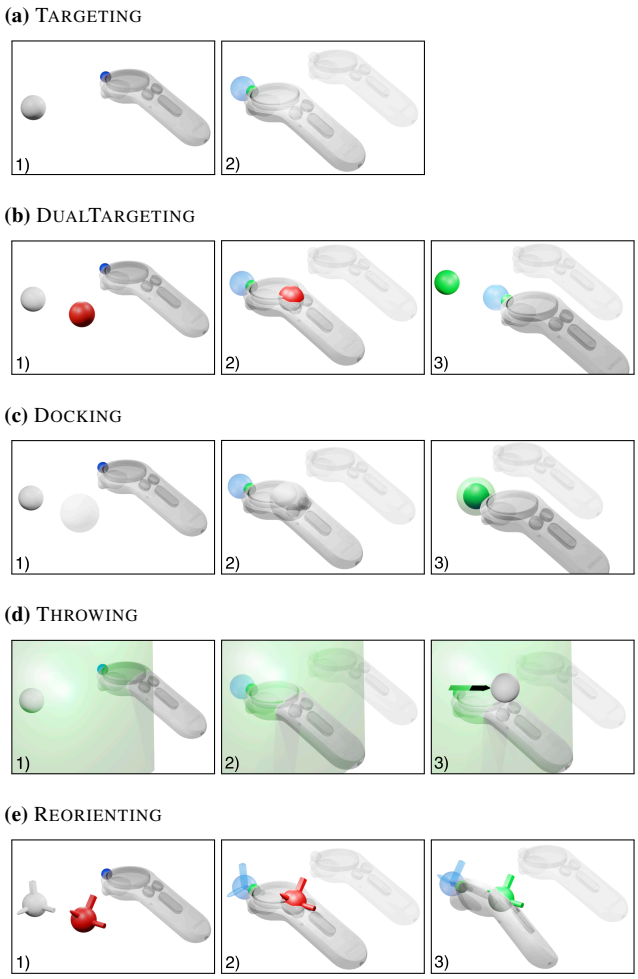

Figure 5: Illustration of visualization, selection and manipulation actions in each task. 1) ***Primary Target (white) appears at the left of the vanished Start Target and Secondary Target appears at the back of the Primary Target. 2) Correct selection on the Primary Target. 3) Perform the manipulation actions on either Primary Target or Secondary Target based on the task.***

for the participant. The experiment consisted of repeated blocks of trials where participants had to select a target and complete a given *Task* on that target using the controller in their dominant hand.

Each trial, therefore, consisted of the following steps. First, the participant had to acquire the *Start Target* to calibrate a starting position. Correctly selecting this target would hide it, start a countdown timer, and reveal a *Primary Target* that the participant was asked to acquire and manipulate to fulfill the *Task* (Section 3.3.1) for the given condition. A trial ended once the *Task* was completed successfully or if the countdown timer exceeded 15 seconds. The countdown timer was implemented to avoid excessive trial completion times, i.e., to limit study duration. The experimental system then moved to the next trial and revealed the *Start Target* again (see Fig. 4 & Fig. 2).

Participants were instructed to complete the trials as quickly and accurately as possible while keeping an error rate below 5% (error rate was displayed to the participant at the end of each block). In case of an error, the trial was appended to the end of the current block, and participants needed to complete all trials within a block successfully before moving on to the next block, thus ensuring the same number of correct trials for all conditions. While participants

were not required to repeat a block if the error rate exceeded 5%, controlling the error rate is a common practice to balance speed and accuracy in pointing experiments [7]. Participants were allowed to take a break without a time limit between each block and after completing each *Task*.

3.5 Data Collection

All participant movement data was logged as a sequence of time-stamped, three-dimensional coordinates. The system also logged selection actions, whether the selection was an error and subsequent task manipulations.

Recall that the goal of this experiment is to measure the effect of different *Task*s (intended uses of the *Primary Target*) on the time taken to select the *Primary Target*. The dependent variables collected and logged by the system were selection time (T_{sel}) and errors. For all trials, we use button press down events on the *Start Target* and *Primary Target* to determine the beginning and end of the *Primary Target*'s selection movement. The selection time was the time interval from the button press action of the *Start Target* to the button press action on the *Primary Target*. The length of time to complete the *Task* was logged, but is immaterial to this experiment as we are only interested in the impact of *Task* on *Primary Target* selection time.

As noted above, the system also identified errors. We consider errors that occurred only while acquiring the *Primary Target*. Errors were classified as one of three error types:

- enter error — when participants pressed the button outside the *Primary Target* before having entered it.
- exit error — when participants pressed the button outside the *Primary Target* after having entered it at least once
- pass error — when the countdown timer reached 0 before the *Primary Target* was correctly selected.

The Error rate, displayed to participants and in our analysis, refers to the percentage of erroneous trials in a block over the total number of trials in the block (including any repeated trials). For example, if a participant failed once on every trial within a block and then succeeded on the second attempt, the error rate would be 50%.

3.6 Design Summary

We adopted a repeated-measures within-subjects design. We effectively looked at three independent variables (IVs): *Task* (TARGETING, DUALTARGETING, DOCKING, THROWING and REORIENTING), *ID* (1.49, 1.81, 2.00, 2.07, 2.37, and 2.66 bits), and *Block* (1-4). The order of *Task* was counterbalanced across participants using a Latin square [58]. Note that for each *ID*, *Primary Direction* **D** and *Secondary Direction* $\mathbf{D}'$ were randomly ordered for generalization. The combination between **D** and *ID* and that of **D** and $\mathbf{D}'$ were not controlled in our experiment. In summary, each participant completed 4 *Block*s × 6 *ID*s × 6 *Primary Direction*s × 5 (counterbalanced) *Task*s, i.e. 720 trials. This resulted in a data set containing 10 800 successful trials for our 15 participants.

4 RESULTS

We analyze our results in terms of error rate, *Primary Target* selection time (T_{sel}), motion kinematics and Fitts' Law modelling. Note that our focus is on T_{sel}, the time taken to move from the *Start Target* to the *Primary Target* given that our interest is on how *Task*s performed on the *Primary Target* impact the time to select it, T_{sel}.

4.1 Error Rate

Recall that, in order to complete a block, participants were required to successfully complete each trial. In the case of an error in selecting the *Primary Target*, participants would need to repeat the corresponding trial at the end of the Block. As a result, alongside the 10,800 correct trials, we collected an additional 289 erroneous trials, for a total of 11,089 trials. Among 289 trials, only 3 trials had pass error and participants did not report the time pressure during the study, implying that the timer did not push participants.

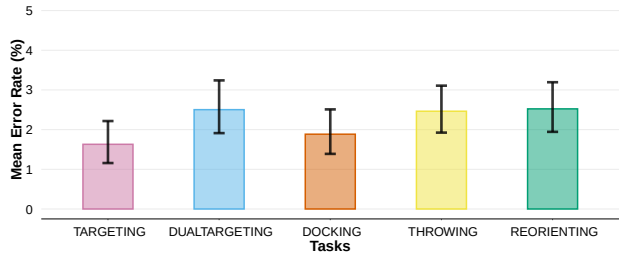

Figure 6: Mean error rate of T_{sel} for *Task*. Error bars are shown with 95% confidence intervals.

Given the non-normal distribution of error rate, a Friedman test was conducted for three independent variables (IVs): *Task*, *Block*, and *ID*. We found a significant effect of *Block* on error rate ($\chi^2_{Block}(3)$=15.46, p<0.005). However, pairwise Wilcoxon rank sum tests with Bonferroni corrections did not show significant differences between blocks. With all blocks, error rate (M=2.20%, SD=5.77) was below the 5% error rate threshold that we recommended our participants not to exceed. Error rate of each *Task* was shown in Fig. 6. The Friedman test did not reveal a significant effect of *Task* on error rate, but showed a significant effect of *ID* ($\chi^2_{Block}(3)$=12.92, p<0.05). However, pairwise comparisons did not reveal any significant differences across *ID*s.

4.2 Selection Time

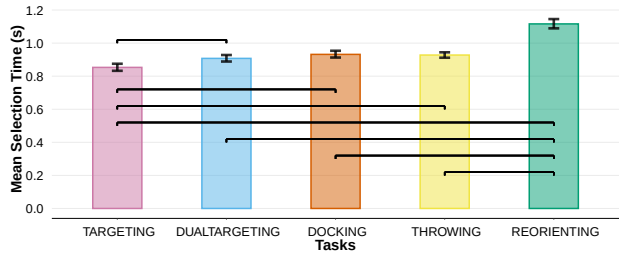

Figure 7: Mean selection time of T_{sel} for *Task*. Error bars are shown with 95% confidence intervals. The statistic significances evaluated by pairwise t-test are connected with lines (p<0.01).

We aggregated non-erroneous trials and removed outliers by eliminating any trial whose selection time was more than three standard deviations from the mean, leaving 10,645 trials for analysis.

Given the non-normal distribution of the data, a Box-Cox transformation [10] was applied to selection time. When sphericity was violated using Mauchly's test, Greenhouse-Geisser correction to the DoFs was applied. When significant effects were found, pairwise t-tests with Bonferroni corrections were conducted for post-hoc analysis. Effect sizes were reported as partial eta squared (η_p^2).

We first conducted a two-way RM-ANOVA (α=0.05) for selection time on *Block* and *Task* to test for a possible learning effect. We found a significant effect of *Block* ($F_{3,42}$=41.64, p<0.001, η_p^2=0.75). Pairwise comparisons revealed significant differences between Block 1 (M=1.03s) and the other three blocks (p<0.001, 2: 0.97s, 3: 0.94s & 4: 0.93s). We found a significant effect of *Task* ($F_{4,56}$=33.11, p<0.001, η_p^2=0.70) but we did not find a significant interaction effect between *Block* and *Task*. These results suggested a potential learning effect that did

not differ between tasks; we thus removed the first block in our remaining analysis.

After removal of the first *Block*,we found a significant effect of *Task* ($F_{4,56}$=29.11, p<0.001, η_p^2=0.68) on T_{sel}. As shown in Fig. 7, pairwise comparisons showed that TARGETING (0.85*s*) was significantly faster than the other four tasks: p<0.001 for all, DUALTARGETING (0.91*s*), DOCKING (0.93*s*), THROWING (0.93*s*) and REORIENTING (1.12*s*). Moreover, REORIENTING was found significantly slower than other tasks: p<0.001 for all. We found a significant effect of *ID* ($F_{2.1,29.7}$=254.54, p<0.001, η_p^2=0.95). Pairwise comparisons showed significant differences (p<0.01) between each ID except for 1.81&2 bits and 2.07&2.37 bits. We also found a significant interaction effect between *Task* and *ID* ($F_{20,280}$=3.00, p<0.001, η_p^2=0.18). Pairwise comparisons showed that there were no significant differences between 1.49&1.81 bits for all tasks, 1.49&2 for DOCKING, DUALTARGETING and REORIENTING, and 2&2.07 for REORIENTING, but that all other levels of *ID* were significantly different for all tasks (p<0.05).

4.3 Motion Kinematics

In order to try to better understand where the difference for T_{sel} between *Tasks* comes from, we analyzed the motion kinematics profile of participants' correct target selections. We once again kept all non-erroneous trials and removed the first block due to the aforementioned possible learning effect.

We computed velocity profiles based on the cursor's 3D position and corresponding timestamps. A time interval of each trial was normalized. Then, the corresponding velocities and normalized distances from the cursor to the *Primary Target* were interpolated to create a time-equidistant profile (every 2% of trial time). Next, for each task and participant, these values were aggregated by computing the average for each normalized time interval. These were subsequently averaged over all participants to produce a single normalized profile for each task (see Fig. 9).

In order to compare our results to [37], we also used these profiles to compute the following motion kinematics measures: peak velocity (pkV), time to peak velocity (t2pkV), percent of time after peak velocity (afterpkV%), and selection velocity (sV: the velocity at the end of the T_{sel}) at *Primary Target* for each trial (see Fig. 8).

For each metric, we removed outliers if the observed value was more than three standard deviations from the mean. Given the non-normal distribution of dependent variables, we conducted an Aligned Rank Transform [60] for pkV, t2pkV, afterpkV% and sV on two IVs: *Task* and *ID*. Contrasts ART [19] with Bonferroni corrections was applied as the post-hoc analysis if a significant effect was found and effect sizes were reported as partial eta squared (η_p^2).

Peak Velocity (Fig. 8 (A)): We found a significant effect of *Task* on pkV ($F_{4,1306}$=76.67, p<0.001, η_p^2=0.19). Contrasts ART showed significant differences of pkV (p<0.001 for all) between each *Task* except between DUALTARGETING (*M*=51.52 cm/s) and DOCKING (51.52 cm/s). We found a significant effect of *ID* on pkV ($F_{5,1306}$=608.60, p<0.001, η_p^2=0.70). Pairwise comparisons revealed significant differences (p< 0.001) between pairs of *ID* except for 1.49&2, 1.81&2.37 and 2.07&2.66 bits. We also found an interaction effect between *Task* and *ID* ($F_{20,1306}$=2.14, p<0.005, η_p^2=0.03). Contrasts ART also showed that there were no significant differences between 1.49&2.00, 1.81&2.37 and 2.07&2.66 bits for each task. In other words, no significant differences on pkV were revealed when target amplitudes were the same.

Time to Peak Velocity (Fig. 8 (B)): We found a significant effect of *Task* on t2pkV ($F_{4,1306}$=43.31, p<0.001, η_p^2=0.12). Pairwise comparisons showed that TARGETING (0.51 s) and DUALTARGETING (0.52 s) had significantly shorter t2pkV (p<0.001 for all) than the other three tasks. We found a significant effect of *ID* ($F_{4,1306}$=27.56, p<0.001, η_p^2=0.10). Pairwise comparisons showed significant differences (p< 0.001) between each *ID* except for 1.49&2.00, 1.81&2.37, 2.07&2.37 and 2.37&2.66 bits. We did not find an interaction effect between *Task* and *ID*.

Percent of Time after Peak Velocity (Fig. 8 (C)): A significant effect of *Task* was revealed on afterpkV% ($F_{4,1306}$=182.49, p<0.001, η_p^2=0.36). The post-hoc analysis showed that TARGETING (38.74%) had significantly smaller afterpkV% (p<0.001 for all) than the other four tasks while REORIENTING (49.53%) had significantly larger afterpkV% (p<0.001 for all). We found a significant effect of *ID* on afterpkV% ($F_{4,1306}$=200.95, p<0.001, η_p^2=0.43). Pairwise comparisons showed significant differences (p< 0.001) between each *ID* except for 2.00&2.07 bits. We did not find an interaction effect between *Task* and *ID*.

Selection Velocity (Fig. 8 (D)): We found a significant effect of *Task* on sV ($F_{4,1306}$=123.43, p<0.001, η_p^2=0.27) and the contrasts ART revealed that REORIENTING had significantly slower sV (7.70 cm/s, p<0.001 for all) and THROWING had significantly higher sV (11.87 cm/s, p<0.001 for all) than other tasks. DUALTARGETING (10.59 cm/s) also had significantly higher sV than TARGETING (9.64 cm/s, p<0.01). We found a significant effect of *ID* on sV ($F_{5,1306}$=95.62, p<0.001, η_p^2=0.27). Pairwise comparisons showed significant differences (p< 0.01) between each *ID* except for 1.49&1.81, 1.49&2.07, 1.81&2.07, 2&2.37 and 2.37&2.66 bits. We also found an interaction effect between *Task* and *ID* ($F_{20,1306}$=2.51, p<0.001, η_p^2=0.04). Pairwise comparisons showed that there were no significant differences between 1.49&1.81, 1.49&2.07, 1.81&2.07, 2.00&2.37, 2.37&2.66 and 2.00&2.66 bits in each task. Interpreting these numbers in terms of target width and amplitude, these results argue that sV is significantly impacted by target width.

4.4 Fitts' Law

Given our research question, i.e. whether *Task* users will perform on an object impacts the time taken to acquire the object, Fitts' law modelling and throughput analysis were applied only to the selection of *Primary Target*, i.e., only to T_{sel} in Figure 2 but not to the subsequent manipulation task. To perform our analysis, we computed the effective target width (W_e) for each target width by multiplying the standard deviation by 4.133 [36, 40] and used W_e to calculate the effective ID, ID_e accordingly. As a result, the difference between W and W_e ranges from 0 cm to 0.6 cm and the corresponding difference between *ID* and ID_e ranges from 0 bits to 0.12 bits, as shown in Table 1.

The classical Fitts' Law was used for modelling [21, 35] due to two concerns: 1. for ID, as our targets are spheres, Fitts' law variants for targets in arbitrary shapes are not necessary. 2. as [54] points out, no work has included all spatial factors in 3D space and a standard metric for 3D modelling is missing. Meanwhile, the classical formula is still a common practice in VR [8, 53, 54].

A(cm)	W(cm)	W_e(cm)	ID(bits)	ID_e(bits)
9.0	3.0	3.2	2.00	1.93
9.0	5.0	4.4	1.49	1.61
12.5	3.0	3.0	2.37	2.37
12.5	5.0	4.6	1.81	1.89
16.0	3.0	3.0	2.66	2.66
16.0	5.0	4.8	2.07	2.12

Table 1: Effective target width (W_e) and effective ID (ID_e) are calculated for each pair of target amplitude (A) and target width (W).

To compensate for the non-normal distribution of selection time in T_{sel}, we adopted a common practice [13, 51] where we first aggregated the mean for each effective Index of Difficulty (ID_e) per *Block* per *Task* and then aggregated the median for each ID_e and *Task*. ID_e ranged from 1.61 to 2.66 bits and the aggregated median time of all tasks correlate with ID_e positively ($R^2 \geq 0.96$), as shown in Fig. 10. When looking at the coefficients of these linear regression models in Table 2, we noticed that slope values b were relatively similar across *Task* but REORIENTING has a higher y-intercept value a than the

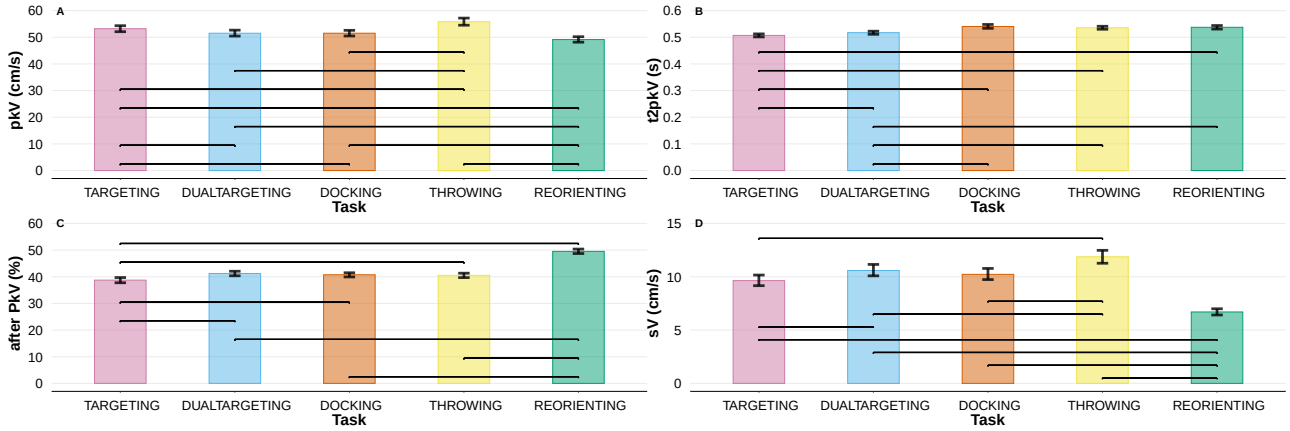

Figure 8: Mean pkV (A), t2pkV (B), afterpkV% (C) and sV (D) of T_{sel} for *Task*. Error bars are shown with 95% confidence intervals. The statistic significances evaluated by ART are connected with lines (p<0.01).

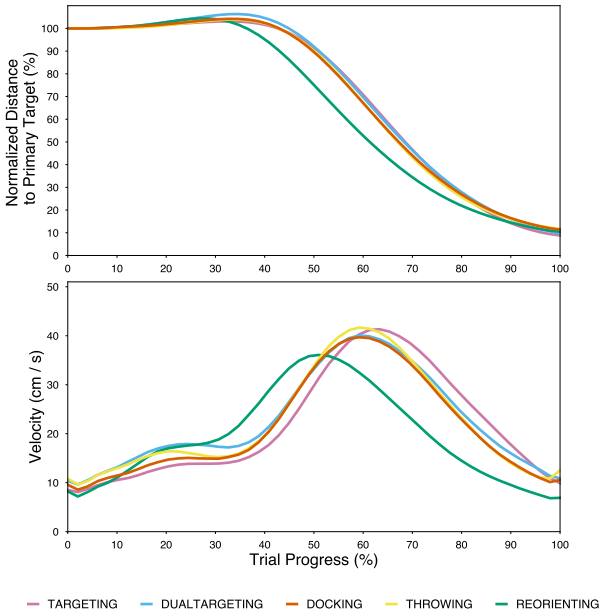

Figure 9: Velocity profile and normalized distance percentage to *Primary Target* in T_{sel}. Trial progress refers to the normalized time interval of a trial.

Task	a	b	TP
TARGETING	0.36 [0.28 0.43]	0.23 [0.20 0.28]	2.47 [2.34 2.61]
DUALTARGETING	0.44 [0.32 0.57]	0.22 [0.16 0.28]	2.29 [2.15 2.44]
DOCKING	0.39 [0.30 0.49]	0.25 [0.21 0.30]	2.25 [2.13 2.38]
THROWING	0.37 [0.27 0.48]	0.26 [0.21 0.31]	2.26 [2.13 2.38]
REORIENTING	0.54 [0.38 0.70]	0.27 [0.20 0.35]	1.87 [1.76 1.99]

Table 2: Modelling results of Fitts' Law ($MT = a + b \cdot ID_e$) and throughput values $TP = ID_e/MT$: estimates (95% CI)

other four tasks. We also report throughput scores in Table 2. Unsurprisingly, TARGETING had the largest throughput, REORIENTING had the lowest throughput, and the throughputs of the other three tasks were similar, a result consistent with T_{sel} values.

5 General Discussion

The goal of this paper is to explore the impact of intended use of a target, i.e. the *Task* performed on the target, on the time required to select that target, i.e. T_{sel}, before performing the manipulation. We examined five tasks: classical TARGETING, DUALTARGETING, DOCKING, THROWING, and REORIENTING. Our hypothesis is that there is an impact, i.e. that the selection that precedes target manipulation is impacted by the specific manipulation *Task* that we perform on the selected target.

Our results support this hypothesis. While we did not find a significant effect on error rate, we did find significant differences in the time taken to select the target. In particular, we found that selection preceding classical TARGETING took significantly less time than selections preceding all other *Task*s we tested, and that selection preceding a REORIENTING *Task* took significantly longer than selections preceding all other tasks. Target selection preceding DUALTARGETING, DOCKING and THROWING did not differ significantly in the time taken. Given that there are significant differences for TARGETING and REORIENTING, we reject our null hypothesis (no impact of intended use) and claim evidence for the impact of intended use on target selection.

To understand where and how this difference occurred, we analyzed movement time during the selection task. First, we created kinematic profiles of distance and speed during target selection. A visual inspection of Figure 9 shows that REORIENTING has a lower peak speed and reaches peak speed earlier in movement with a corresponding increase in the length of time spent decelerating. Further kinematic analysis supports this visual analysis; in Figure 8, REORIENTING has the lowest peak speed (pkV) and the longest deceleration phase (greater $afterpkV\%$) than other *Task*s by a significant margin. It implied that participants planned their movement and did not rush during the selection for REORIENTING. Besides, the Fitts' law modelling confirmed this by showing that the reaction time for REORIENTING was higher and the throughput value of REORIENTING was smaller than other tasks.

The kinematic analysis does, however, present some additional observations that merit future investigation. For example, while REORIENTING results in the lowest peak speed during the preceding selection, THROWING results in the preceding selection movement characterized by the highest peak speed, and this peak speed differs

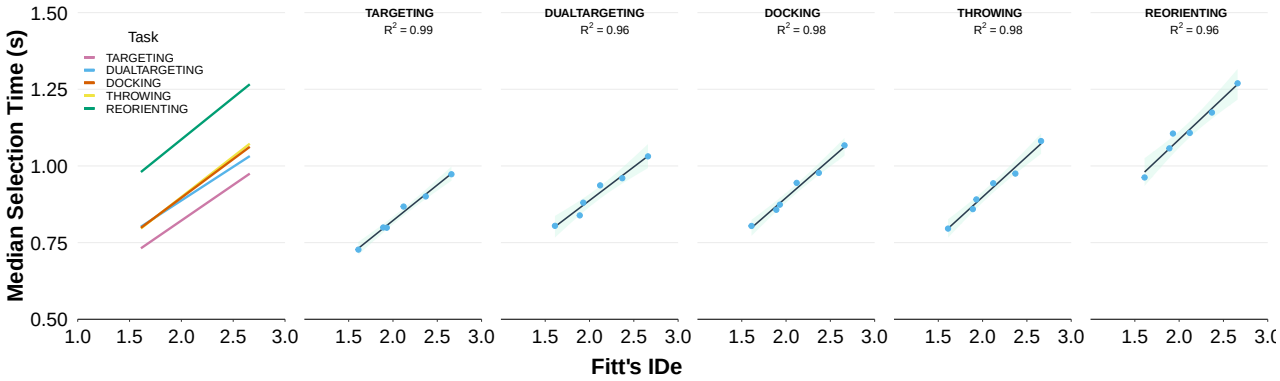

Figure 10: Median Selection Time as a function of Fitts' ID_e per task, with corresponding R^2 and 95% confidence interval (light-green area).

significantly from all other peak speeds. This is followed by peak speed for selection preceding TARGETING. DUALTARGETING and DOCKING did not result in significant differences in peak speed in the preceding selection kinematics. Effect sizes are not small for various measures of kinematics highlighted in Section 4.3. Based on typical interpretations of effect size, $\eta^2 > 0.14$, i.e. large effect size[1], for all differences except time to peak speed ($t2pkV$), where the effect size maps to medium effect. The absence of additional significant differences in selection time between DUALTARGETING, DOCKING and THROWING does not imply that differences in selection prior to these *Tasks* do not exist; it simply implies that we measured no significant differences in selection time. Future work is planned to probe these effects in more depth.

It is interesting to contrast our results in Mandryk's and Lough's results in 2D [37]. Mandryk's and Lough's four intended uses (*Tasks*) correspond to TARGETING, DUALTARGETING, DOCKING, and THROWING in our experiment. They found that TARGETING and DUALTARGETING resulted in selection times (T_{sel}) that differed significantly from DOCKING's and THROWING's selection times, but did not observe differences between TARGETING and DUALTARGETING. In contrast, while we, too, found that TARGETING resulted in shortest preceding selection times, DUALTARGETING resulted in T_{sel} more similar to DOCKING's and THROWING's T_{sel}. In terms of overall differences, we and Mandryk and Lough find differences in T_{sel} of just under 10% in overall magnitude for these selections (9.4% in ours and 8.8% estimated in theirs). Our REORIENTING is unique to our study, and resulted in the most significant differences in preceding selection time; REORIENTING resulted in target selection times more than 30% longer (1.12s vs 0.85s) than TARGETING.

Given our results that the intended use of a target impacts the time taken to select a target, the question that follows from this is what these results mean. To address this question, we point, again, to the analysis of 20 years of VR-based by Bergström et al. [9]. As noted earlier in our review of this work, Bergström et al. highlight that, as dependent variables, selection task studies measure selection time for the selection task and manipulation task studies measure the time participants take to manipulate an object. Success in selection-based research is measured by shorting selection times or reducing errors, or both; success in manipulation is similarly based on increased throughput for the manipulation task. The assumption that underlies these success metrics is, *ipso facto*, that each individual user action can be optimized in isolation from other tasks in interfaces, i.e., that manipulation does not impinge upon preceding selection, but our results argue that this assumption cannot be made. If different manipulations impact the time taken for a preceding selection, then measuring only the manipulation time may over-estimate (or under-estimate) the benefits of a novel interaction technique.

It is also true that we only measure retrospective impact (i.e. the impact of future intended use on preceding selection), but prospective impacts are also possible. It is hypothetically possible that a pointing facilitation technique in VR, e.g. an area cursor, might impact a user's ability to perform a task on a target, e.g. reorientation of the target acquired via the area cursor. While exploring prospective impacts is one area of future work, it also highlights a more general implication for system design. Specifically, when we have available to us new interaction techniques (i.e. new manipulations) or new pointing facilitation techniques (i.e. new selection techniques), incorporating them into realistic systems requires thinking not just about the individual action that they optimize but also about their place within and more general impacts on the overall task flow of the user.

5.1 Applicable Scenarios

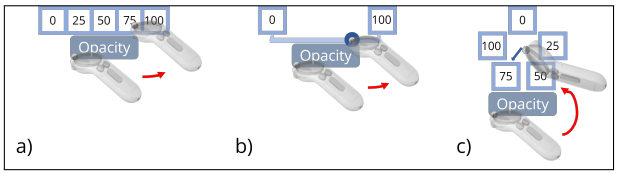

Figure 11: Menu selection for changing target opacity: a) DUAL-TARGETING, b) DOCKING, c) REORIENTING.

Our results can be framed into concrete applications and interface design in VR. One classical example is menu selection [20, 39]. Considering Fig. 11, our results suggest that interactions that take advantage of DUALTARGETING and DOCKING may result in similar selection time prior to the manipulation sub-task and differ from each other based on the design of manipulation techniques. In contrast, techniques that leverage REORIENTING already take longer selection time prior to the manipulation. Similarly, in data visualization tasks in virtual environments [11], rotational manipulations of a dataset may introduce additional costs if the target acquisition prior to the rotational manipulation is slowed.

5.2 Limitations

One limitation that we highlight is that our experiment was conducted remotely in participants' homes. While an in-home environment increases the external validity of our study, it cannot be as controlled as a laboratory one, whose setup and control can assume to be optimized for the experiment. Space in homes may be constrained, and households may present interruptions during the experiment. To limit this as a factor, we note that participants

[1] https://www.spss-tutorials.com/effect-size/

were provided with detailed instructions, from both pre-recorded videos and experiment systems. A researcher was also present via video conferencing tools during the experiment, allowing the environment to be monitored for confounds. We also note that our within-subjects design partially controls for confounds by ensuring that the environment is similar across each task for a participant.

While we contrast five different *Task*s (the intended use) in our study, there exist more complex manipulations in the virtual environments. Furthermore, we do not consider objects' surface characteristics, perceived weight, and perceived fragility [29, 46] as objects in our study have similar surface, weight and fragility. We also do no consider bimanual interactions [26] nor co-articulated actions (e.g. 6-dof reorienting, i.e. a docking task that requires both rotation and translation) [28, 48].As noted in the experiment design, we did not control the *Primary Direction* and *Secondary Direction*. Further research can explore how directions can impact the acquisition for various intended uses, particularly for co-articulated actions. Finally, we restrict our current experiment to arms' length interactions with a controller on static targets, so target widths and distances are constrained, i.e., we have IDs between 1.49 bits to 2.66 bits. While these IDs are commonly used in a arms' length interactions in VR environment, they are low compared to desktop interfaces [37], touch-based interfaces [23], distant interactions in VR [7, 8, 13, 14], contexts where movement amplitudes can increase due to the greater distance of targets from the user. This also explains why the throughput values in Table 2 are relatively lower, compared to throughput scores in other VR studies [8, 53]. Incorporating more distant interactions does, however, add additional complexity to the selection action because direction (e.g. targeting via a ray) and depth are often controlled differently during distant interaction [7, 8].

5.3 Future Work

The results of our intended use study present interesting avenues of future research into interaction in virtual reality environments. As one example, a novel interaction technique might not result in higher throughput during the interaction, but it is possible that it might speed the selection that precedes the interaction. In this case, considering both the selection and manipulation of a target as a unified task allows us to identify potentially beneficial novel interactions that might have been ignored if the only metric for success is throughput for the manipulation.

Another possible area of inquiry given differences in kinematics noted in Section 4.3 is that a more careful analysis of movement during selection might allow the system to infer what a user intends to do with a target prior to acquiring the target. This, in turn, could allow us to develop interaction techniques that leverage this inference. Reorientation appears a good initial candidate to identify given the deviation in the selection kinematics shown in Figure 9.

6 Conclusion

A significant body of research in virtual reality exists that explores both selection and manipulations techniques to speed performance, and this past research typically assumes that these actions – the selection and the subsequent manipulation – can be independently optimized [9]. Based upon past work in 2D environments, in this paper we question this assumption. We examine the impact of five common virtual reality manipulations (Targeting, DualTargeting, Docking, Throwing, and Reorienting) on the time taken to select a target prior to performing the manipulation. We identify the existence of an effect of subsequent use of a target on the act of selecting the target. Specifically, we find that Targeting had the shortest selection time, Reorienting had the longest, and the other three intended uses we evaluate result in acquisition times between these two extremes. We synthesize these results and highlight their implications for research and design in VR environments.

References

[1] J. Accot and S. Zhai. Beyond fitts' law: Models for trajectory-based hci tasks. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*, CHI '97, p. 295–302. Association for Computing Machinery, New York, NY, USA, 1997. doi: 10.1145/258549.258760

[2] J. Accot and S. Zhai. Refining fitts' law models for bivariate pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, p. 193–200. Association for Computing Machinery, New York, NY, USA, 2003. doi: 10.1145/642611.642646

[3] Half-life: Alyx, 2022.

[4] R. Arora, R. H. Kazi, D. M. Kaufman, W. Li, and K. Singh. Magicalhands: Mid-air hand gestures for animating in vr. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, UIST '19, p. 463–477. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3332165.3347942

[5] R. Arora, J. Li, G. Shi, and K. Singh. Thinking outside the lab: VR size & depth perception in the wild, 2021.

[6] R. Balakrishnan. "beating" fitts' law: virtual enhancements for pointing facilitation. *International Journal of Human-Computer Studies*, 61(6):857–874, 2004. Fitts' law 50 years later: applications and contributions from human-computer interaction. doi: 10.1016/j.ijhcs.2004.09.002

[7] M. Baloup, T. Pietrzak, and G. Casiez. *RayCursor: A 3D Pointing Facilitation Technique Based on Raycasting*, p. 1–12. Association for Computing Machinery, New York, NY, USA, 2019.

[8] A. U. Batmaz and W. Stuerzlinger. Effective throughput analysis of different task execution strategies for mid-air fitts' tasks in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–9, 2022. doi: 10.1109/TVCG.2022.3203105

[9] J. Bergström, T.-S. Dalsgaard, J. Alexander, and K. Hornbæk. *How to Evaluate Object Selection and Manipulation in VR? Guidelines from 20 Years of Studies*. Association for Computing Machinery, New York, NY, USA, 2021.

[10] G. E. P. Box and D. R. Cox. An analysis of transformations. *Journal of the Royal Statistical Society. Series B (Methodological*, pp. 211–252, 1964.

[11] W. Büschel, A. Mitschick, T. Meyer, and R. Dachselt. Investigating smartphone-based pan and zoom in 3d data spaces in augmented reality. In *Proceedings of the 21st International Conference on Human-Computer Interaction with Mobile Devices and Services*, MobileHCI '19. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3338286.3340113

[12] S. K. Card, T. P. Moran, and A. Newell. The keystroke-level model for user performance time with interactive systems. *Commun. ACM*, 23(7):396–410, July 1980. doi: 10.1145/358886.358895

[13] Y. Chen, K. Katsuragawa, and E. Lank. Understanding viewport- and world-based pointing with everyday smart devices in immersive augmented reality. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, p. 1–13. Association for Computing Machinery, New York, NY, USA, 2020. doi: 10.1145/3313831.3376592

[14] Y. Chen, J. Sun, Q. Xu, E. Lank, P. Irani, and W. Li. Empirical evaluation of moving target selection in virtual reality using egocentric metaphors. In C. Ardito, R. Lanzilotti, A. Malizia, H. Petrie, A. Piccinno, G. Desolda, and K. Inkpen, eds., *Human-Computer Interaction – INTERACT 2021*, pp. 29–50. Springer International Publishing, Cham, 2021.

[15] Y. Chen, J. Sun, Q. Xu, E. Lank, P. Irani, and W. Li. Global scene filtering, exploration, and pointing in occluded virtual space. In *IFIP Conference on Human-Computer Interaction*, pp. 156–176. Springer, 2021.

[16] A. Cockburn and B. McKenzie. Evaluating the effectiveness of spatial memory in 2d and 3d physical and virtual environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '02, p. 203–210. Association for Computing Machinery, New York, NY, USA, 2002. doi: 10.1145/503376.503413

[17] C. Donalek, S. G. Djorgovski, A. Cioc, A. Wang, J. Zhang, E. Lawler, S. Yeh, A. Mahabal, M. Graham, A. Drake, et al. Immersive and

collaborative data visualization using virtual reality platforms. In *2014 IEEE International Conference on Big Data (Big Data)*, pp. 609–614. IEEE, 2014.

[18] K. El Batran and M. D. Dunlop. Enhancing klm (keystroke-level model) to fit touch screen mobile devices. In *Proceedings of the 16th International Conference on Human-Computer Interaction with Mobile Devices & Services*, MobileHCI '14, p. 283–286. Association for Computing Machinery, New York, NY, USA, 2014. doi: 10.1145/2628363.2628385

[19] L. A. Elkin, M. Kay, J. J. Higgins, and J. O. Wobbrock. An aligned rank transform procedure for multifactor contrast tests. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, p. 754–768. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3472749.3474784

[20] K. Fennedy, J. Hartmann, Q. Roy, S. T. Perrault, and D. Vogel. Octopocus in vr: Using a dynamic guide for 3d mid-air gestures in virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2021. doi: 10.1109/TVCG.2021.3101854

[21] P. M. Fitts. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*, 47 6:381–91, 1954.

[22] M. Gentilucci. Object motor representation and reaching–grasping control. *Neuropsychologia*, 40(8):1139–1153, 2002. doi: 10.1016/S0028-3932(01)00233-0

[23] A. Goguey, M. Nancel, G. Casiez, and D. Vogel. The performance and preference of different fingers and chords for pointing, dragging, and object transformation. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, p. 4250–4261. Association for Computing Machinery, New York, NY, USA, 2016. doi: 10.1145/2858036.2858194

[24] Google earth vr, 2022.

[25] E. D. Graham and C. L. MacKenzie. Physical versus virtual pointing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '96, p. 292–299. Association for Computing Machinery, New York, NY, USA, 1996. doi: 10.1145/238386.238532

[26] Y. Guiard. Asymmetric division of labor in human skilled bimanual action. *Journal of Motor Behavior*, 19(4):486–517, 1987. PMID: 15136274. doi: 10.1080/00222895.1987.10735426

[27] E. R. HOFFMANN. A comparison of hand and foot movement times. *Ergonomics*, 34(4):397–406, 1991. PMID: 1860460. doi: 10.1080/00140139108967324

[28] T. E. Jerde, J. F. Soechting, and M. Flanders. Coarticulation in fluent fingerspelling. *Journal of Neuroscience*, 23(6):2383–2393, 2003. doi: 10.1523/JNEUROSCI.23-06-02383.2003

[29] R. S. Johansson and J. R. Flanagan. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–359, 2009.

[30] N. Klopcar and J. Lenarcic. Kinematic model for determination of human arm reachable workspace. *Meccanica*, 40:203–219, 2005.

[31] P. Knierim, V. Schwind, A. M. Feit, F. Nieuwenhuizen, and N. Henze. Physical keyboards in virtual reality: Analysis of typing performance and effects of avatar hands. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, p. 1–9. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3173574.3173919

[32] G. Kurillo, A. Chen, R. Bajcsy, and J. Han. Evaluation of upper extremity reachable workspace using kinect camera. *Technology and health care : official journal of the European Society for Engineering and Medicine*, 21, 11 2013. doi: 10.3233/THC-130764

[33] E. Lank, Y.-C. N. Cheng, and J. Ruiz. Endpoint prediction using motion kinematics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, p. 637–646. Association for Computing Machinery, New York, NY, USA, 2007. doi: 10.1145/1240624.1240724

[34] L. Li, F. Yu, D. Shi, J. Shi, Z. Tian, J. Yang, X. Wang, and Q. Jiang. Application of virtual reality technology in clinical medicine. *American journal of translational research*, 9(9):3867, 2017.

[35] I. S. MacKenzie. A note on the information-theoretic basis for fitts' law. *Journal of Motor Behavior*, 21(3):323–330, 1989. PMID: 15136269. doi: 10.1080/00222895.1989.10735486

[36] I. S. MacKenzie. Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction*, 7(1):91–139, 1992. doi: 10.1207/s15327051hci0701$_3

[37] R. L. Mandryk and C. Lough. *The Effects of Intended Use on Target Acquisition*, p. 1649–1652. Association for Computing Machinery, New York, NY, USA, 2011.

[38] Meetinvr - business meetings & collaboration in vr, 2022.

[39] M. Mundt and T. Mathew. *An Evaluation of Pie Menus for System Control in Virtual Reality*. Association for Computing Machinery, New York, NY, USA, 2020.

[40] A. Murata. Extending effective target width in fitts' law to a two-dimensional pointing task. *International journal of human-computer interaction*, 11(2):137–152, 1999.

[41] M. Nancel and E. Lank. *Modeling User Performance on Curved Constrained Paths*, p. 244–254. Association for Computing Machinery, New York, NY, USA, 2017.

[42] D. Norman. Affordance, conventions, and design. *Interactions*, 6:38–42, 05 1999. doi: 10.1145/301153.301168

[43] S. Pheasant and C. Haslegrave. *Bodyspace: Anthropometry, Ergonomics and the Design of Work*. 12 2018. doi: 10.1201/9781315375212

[44] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa. The go-go interaction technique: Non-linear mapping for direct manipulation in vr. In *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*, UIST '96, p. 79–80. Association for Computing Machinery, New York, NY, USA, 1996. doi: 10.1145/237091.237102

[45] I. Poupyrev, T. Ichikawa, S. Weghorst, and M. Billinghurst. Egocentric object manipulation in virtual environments: Empirical evaluation of interaction techniques. *Computer Graphics Forum*, 17(3):41–52, 1998. doi: 10.1111/1467-8659.00252

[46] D. A. Rosenbaum. *Human motor control*. Academic press, 2009.

[47] J. Ruiz and E. Lank. Analyzing intended use effects in target acquisition. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces*, AVI '14, pp. 145–152. ACM, New York, NY, USA, 2014. doi: 10.1145/2598153.2598158

[48] D. E. Rumelhart and D. A. Norman. Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive science*, 6(1):1–36, 1982.

[49] C. Seegelke, C. M. Hughes, and T. Schack. An investigation into manual asymmetries in grasp behavior and kinematics during an object manipulation task. *Experimental Brain Research*, 215(1):65–75, 2011.

[50] W. J. Shelstad, D. C. Smith, and B. S. Chaparro. Gaming on the rift: How virtual reality affects game user satisfaction. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 61(1):2072–2076, 2017. doi: 10.1177/1541931213602001

[51] S. Siddhpuria, S. Malacria, M. Nancel, and E. Lank. *Pointing at a Distance with Everyday Smart Devices*, p. 1–11. Association for Computing Machinery, New York, NY, USA, 2018.

[52] K. Spiel, O. L. Haimson, and D. Lottridge. How to do better with gender on surveys: A guide for hci researchers. *Interactions*, 26(4):62–65, June 2019. doi: 10.1145/3338283

[53] R. J. Teather and W. Stuerzlinger. Pointing at 3d targets in a stereo head-tracked virtual environment. In *2011 IEEE Symposium on 3D User Interfaces (3DUI)*, pp. 87–94, 2011. doi: 10.1109/3DUI.2011.5759222

[54] E. Triantafyllidis and Z. Li. The challenges in modeling human performance in 3d space with fitts' law. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI EA '21. Association for Computing Machinery, New York, NY, USA, 2021. doi: 10.1145/3411763.3443442

[55] M. Van Veldhuizen and X. Yang. The effect of semi-transparent and interpenetrable hands on object manipulation in virtual reality. *IEEE Access*, 9:17572–17583, 2021. doi: 10.1109/ACCESS.2021.3050881

[56] E. Velloso, J. Alexander, A. Bulling, and H. Gellersen. Interactions under the desk: A characterisation of foot movements for input in a seated position. vol. 9296, pp. 384–401, 09 2015. doi: 10.1007/978-3-319-22701-6_29

[57] O. Špakov and D. Miniotas. Gaze-based selection of standard-size menu items. In *Proceedings of the 7th International Conference on*

Multimodal Interfaces, ICMI '05, p. 124–128. Association for Computing Machinery, New York, NY, USA, 2005. doi: 10.1145/1088463.1088486

[58] E. Williams. Experimental designs balanced for the estimation of residual effects of treatments. 1949. doi: 10.1071/CH9490149

[59] J. O. Wobbrock, E. Cutrell, S. Harada, and I. S. MacKenzie. An error model for pointing based on fitts' law. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, p. 1613–1622. Association for Computing Machinery, New York, NY, USA, 2008. doi: 10.1145/1357054.1357306

[60] J. O. Wobbrock, L. Findlater, D. Gergle, and J. J. Higgins. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 143–146, 2011.

[61] H. Xia, S. Herscher, K. Perlin, and D. Wigdor. Spacetime: Enabling fluid individual and collaborative editing in virtual reality. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, UIST '18, p. 853–866. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3242587.3242597

[62] S. Zhai. Characterizing computer input with fitts' law parameters—the information and non-information aspects of pointing. *International Journal of Human-Computer Studies*, 61(6):791–809, 2004. Fitts' law 50 years later: applications and contributions from human-computer interaction. doi: 10.1016/j.ijhcs.2004.09.006