



TouchGlass: Raycasting from a Glass Surface to Point at Physical Objects in Public Exhibits

Florent Cabric¹(✉), Emmanuel Dubois¹, Pourang Irani^{1,2},
and Marcos Serrano¹

¹ University of Toulouse, Toulouse, France
{florent.cabric, emmanuel.dubois, pourang.irani,
marcos.serrano}@irit.fr

² HCI Lab, University of Manitoba, Winnipeg, Canada
irani@cs.umanitoba.ca

Abstract. Physical objects such as natural items or fine art pieces are often placed behind glass cases to protect them from dust and damage. Generally, interacting with such objects is indirect, based for example on an adjacent touch interface detracting users' attention from the object. In this paper, we explore whether the glass case could be used as an input surface to point and select distant physical objects. With such an approach, the glass case offers a physical delimiter for interaction to avoid unintended activations. We explore this innovative approach through a two steps approach. First, we carry an informative study with 46 participants to validate the most appropriate “walk-up and use” technique. Our results show that using a ray orthogonal to the glass surface is the most natural approach in a public setting. Next, we further explore this orthogonal raycasting technique and conduct a target acquisition experiment to evaluate the impact on target selection performance of the target size, target distance, presence of spatial references and user's head position with regards to the glass case. Results reveal that using the glass as touch surface allows to easily select targets as small as 3 cm up to 35 cm away from the glass. From these results, we provide a set of guidelines to design interactive exhibits using a touch glass case.

Keywords: Touch input · Distant pointing · Transparent touch surface · Absolute pointing · Evaluation · Physical objects

1 Introduction

Physical objects such as natural history items or fine art pieces are often placed behind a glass case in public exhibits to protect them from dust or damage. Most of these exhibits are not interactive, hence the level of information that can be displayed is usually limited to small physical notes. The few interactions in this context are often

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-29387-1_15) contains supplementary material, which is available to authorized users.

© IFIP International Federation for Information Processing 2019

Published by Springer Nature Switzerland AG 2019

D. Lamas et al. (Eds.): INTERACT 2019, LNCS 11748, pp. 249–269, 2019.

https://doi.org/10.1007/978-3-030-29387-1_15

based on a touch input on a separated display [1], which detracts the user from the exhibit artifact.

While these objects are placed behind a glass, few projects have explored how to use such glass cases to interact with objects inside the case. Previous solutions include the use of augmented reality through smartphones or tablets [2]. Those solutions require the user to install a dedicated application. Alternative proposals include the use of mid-air gestures [3], which suffer from discoverability and delimiter problems. Interacting on the glass case instead has several advantages: (1) it offers a physical delimiter for interaction (touching the glass) avoiding unintended activations; (2) it does not visually deviate the user from the physical content displayed inside the glass case; (3) finally, it is technically robust, low-cost and easy to implement on existing glass cases thanks to current touch technologies (e.g. infrared panels).

In this paper, we explore how to select a physical object behind a glass by using 2D touch input on the glass itself (cf. Fig. 1). We first analyze the different “walk-up and use” approaches that users would employ when faced with a tactile glass. To this end, we ran a study aimed at exploring the users’ spontaneous way of pointing at an object behind a glass. The results reveal that the most frequent interaction is similar to the way objects are selected on touch screens (i.e. through raycasting using a 2D ray perpendicular to the glass surface). However, unlike classical direct interaction on touchscreens, this setup introduces a spatial gap between the object to select (showcased in the glass case) and the touch surface (i.e. the glass). This gap may affect the ability of the user to touch the surface on a place that is associated with the target behind the glass.



Fig. 1. A user selecting different physical objects through a glass case

In a second step, we evaluate the impact of the target size, the target distance, the presence of spatial references and the user’s head position with regards to the glass case, on target selection using an orthogonal ray. This experiment reveals that using the glass to point at objects is a viable approach, as it allows to easily select targets as small as 3 cm up to 35 cm away from the glass.

Our contributions are (1) a first study revealing the best “walk-up-and-use” approach to interact with augmented glass cases; (2) a target acquisition controlled

study assessing the performance of pointing on a glass case using an orthogonal ray, and (3) guidelines for designing pointing systems in interactive exhibits using a glass as touch input.

2 Related Work

Our work focuses on the use of an interactive glass case to point at a distant object, i.e. an object out of user's arms reach. Direct and indirect touch input has been widely studied in different contexts (large surfaces [4], single-pointer interaction [5], pen-based input [6], etc.). The originality of our approach is that the object to select is not just behind the tactile surface, but physically distant from it. In such a context, raycasting is the most widely used technique for direct pointing, while using a separated touch surface is the most frequent technique for indirect pointing. We synthesize works related to these two aspects (*direct vs. indirect*) in the following sections and describe concrete examples of distant pointing techniques in public situations.

2.1 Raycasting as a Direct Technique for Pointing at Distant Objects

Most techniques to select objects in space are related to raycasting, which is the standard technique to point and select objects in virtual environments [7–9]. To control the ray direction in the real world, most approaches used either the finger [10, 11] or a handheld object [12–14]. These techniques have been used in volumetric displays [15, 16] and AR systems [17–19], which can display the entire ray. However, using such raycasting techniques in the real world raises two major problems: the parallax effect and the lack of feedback to render either the ray in midair or the resulting pointer/cursor.

Parallax Effect. In the context of touch surfaces, the parallax effect has been defined as “a difference between the perceived location at a particular UI element and users’ actual touch location on the surface” [20]: it is related to the inability for the user’s eye to focus simultaneously at two different depths (called binocular parallax), e.g. the tip of the finger and the distant object. Errors due to this parallax are designated as “parallax errors” by Migge et Kunz [21]. An example of such effect can be found on public interactive displays protected by a thick tactile glass (e.g. ATM): in this context, Khamis et al. [20] studied the gap between the touch and the pointer. They explored the effect of several correcting methods to increase the accuracy of touch, which improved significantly the accuracy of pointing tasks. In a different context, Lee et al. [22] designed a binocular cursor to increase pointing performance. In both cases, the proposed solutions require to display visual feedback to alleviate the parallax effect. However, in our context, the glass case has no display capabilities.

Lack of Feedback. Concerning the difficulty to provide a feedback when using a raycasting technique for pointing at a distant object, previous works applied projection mapping, which draws the ray projection (i.e. the cursor) on the physical surface. For instance, Bandyopadhyay et al. [23] used a lamp as a pointing device to select physical objects. Some researchers applied laser pointers to represent the ray projection [24].

Even though these approaches usually work well with one large object, they can be limited when dealing with small objects, such as insects in a museum, or when there is no projection surface in between two items, i.e. when they are hanging.

Recently, Freeman et al. [25] investigated pointing at levitating small spherical objects with mid-air movement of the index finger. The technique consisted of a 3D raycast, without any visual feedback of the ray (except for the target shaking when pointed at). However, since they investigated a very small distance between finger and target (from 5 cm to 8 cm), their results are difficult to extend to our usage context (i.e. public exhibit), where objects are usually further away from the glass.

2.2 Indirect Touch-Techniques for Pointing at a Distant Object

There are two types of indirect touch pointing techniques: absolute (i.e. the input device sends its position (x, y)) or relative (i.e. the input device sends its movement $(\Delta x, \Delta y)$) [26]. Relative solutions require to permanently visualize the current position of the pointer. In our context displaying a pointer is not viable (as explained in the previous section), so this section focuses on absolute indirect pointing solutions only.

Previous works explored the main factors influencing absolute indirect pointing. A first important factor is the presence of a tactile spatial reference. For instance, landmarks placed on the back-of-device contribute to increasing accuracy [27]. A second factor is the relative size and aspect ratios between the input and the display, which may negatively influence success rate of target selection [28]. Finally, a third important factor is the type of feedback. Previous studies on touch input for large projected displays revealed that, even without feedback, users can select 90% of targets [29]. This means that absolute pointing is still a good candidate when no feedback is provided. However, these three studies focused on interaction with 2D touchscreens while in our case we investigate the use of a transparent glass surface with a large physical gap between the glass and the target.

2.3 Pointing at Items Behind Glass in Public Exhibits

Previous research proposed a variety of approaches to promote a more engaging experience with exhibit items. For instance, [30] proposed using mid-air gestures, detected with a Kinect, to point at an object behind a glass. Plasencia et al. [3] implemented a similar approach, augmenting the glass case with AR capabilities. The system detected the hand position using a Leap Motion and the augmented glass showed the hand reflection to facilitate the object selection. The flashlight metaphor was investigated in public exhibits [31, 32] through a system that augments physical museum pieces using video projection. Pointing requires using a 6 DOF controller or a Leap Motion. The main limitation of these approaches is that they require to augment the glass cases with some projection mapping. Moreover, mid-air gestures suffer from a lack of delimiter, which can lead to false positives. Finally, the need to spatially track hand and fingers can usually only be performed on a reduced volume.

2.4 Synthesis

On the one hand, raycasting is a popular direct pointing technique for reaching distant objects, but can introduce a parallax effect and will suffer from the lack of feedback in the real world. On the other hand, indirect techniques appear to be very effective even without feedback but tend to break the interaction flow as they split the user's attention between different input and output spaces.

In our work, we explore how to combine these two approaches to take advantage of their respective benefits, which results in the concept of *ray casting from a glass surface*. In a public exhibit context, the use of a ray casting technique should make the interaction easy to discover and learn. The glass can be considered as an indirect input surface, which offers the advantage of providing a physical delimiter and should facilitate pointing without feedback.

To our knowledge, this concept of ray casting from a glass surface has not been studied before and constitutes a different interaction setup than those previously cited: the touch surface is transparent and the object to select is at a distance behind it. Therefore, we carried a two steps approach to explore this new environment: first, we led an informative study to find the most natural interaction in a public context; then, we further explored the most used approach in a controlled experiment to assess its performance.

3 Informative Study

While using a glass to select a physical object placed behind it presents several advantages as mentioned earlier, casting a ray only requires to define an origin and a direction or a target. Hence there are different possible ways of casting a ray from the glass surface: by using the finger orientation [10, 15, 25] (i.e. finger oriented raycasting cf. Fig. 2A), by using an eye-finger raycasting [33–35] (also known as image-plane technique cf. Fig. 2B) or by placing the finger in front of the object and using the glass surface as a touchscreen (i.e. raycasting orthogonal to the glass cf. Fig. 2C). However, an appropriate interaction technique for a public context needs to be as self-explanatory [36] as possible, so that passers-by interact with the glass with a minimum of instructions.

To find out the most appropriate interaction, we conducted an informative study in a public environment with passers-by. The goal of this study was to explore how participants would intuitively make use of the glass to point at an object behind it.

3.1 Experimental Protocol

Task and Instructions. Five white physical spheres with a diameter of 2 cm were positioned on cardboard, placed 55 cm behind the glass. Four of them defined the corners of a 33×26 cm side rectangle, while the fifth one was positioned in the middle of the rectangle. Participants received the following instruction: *“Please point at each of the five white spheres, using the glass case in front of you as a support for the interaction”*. Before pointing at a target, participants were requested to orally




Fig. 2. Three types of raycasting performed by participants: (A) Oriented, (B) Eye-Finger, and (C) Orthogonal. Detailed view of the glass and targets (D).

precise which target they were planning to select. Participants had to select all of them in a free order. Participants were not aware of these three categories when taking part into the experiment.

Participants. We recruited 46 participants (39 males and 7 females) aged between 19 and 62 ($M = 33$). All participants were recruited in the hall of the local university as they passed-by: 13 were university members, 26 were computer science students, five were engineers and two were administrative employees.

Apparatus. The glass case was installed during 10 h in a hall of our local university. The glass panel (78.5×44 cm) was placed on a 90 cm height table. Users had to position themselves in the middle of the glass panel and at a distance where all targets were visible and the glass case easily reachable with their arm (cf. Fig. 2).

Collected Data. All the experiment was video-recorded. After all trials were completed, we asked the participants to provide a graphical representation of how their finger was positioned and oriented to perform each selection task. Based on the video recording and the representations provided by the participants, experimenters clustered the results into three categories, corresponding to the possible approaches presented for casting a ray from the glass: (1) oriented raycasting, i.e. using their finger orientation; (2) eye-finger raycasting, i.e. aiming at the object with the finger; and (3) orthogonal raycasting, i.e. casting a ray perpendicular to the glass. The experimenter cross-checked the video and the graphical representation to ensure there were no inconsistencies between the method used and the one reported by the participants.

3.2 Results

Among the 46 participants, 40 of them used some sort of raycasting to select targets as expected, while 4 participants used gestural input and 2 participants misunderstood the instructions (i.e. they tried to reach the objects by avoiding the glass case). Therefore, our analysis of the type of raycasting focuses on the 40 participants who employed some sort of raycasting to point at the objects behind the glass.




Fig. 3. Percentage of participants that employed each type of raycasting: oriented raycasting (left), eye-finger raycasting (center) and orthogonal raycasting (right).

Our results reveal that among those 40 participants, 5 participants (12,5%) cast an oriented ray (i.e. using their finger orientation) to select targets; 5 participants (12,5%) cast an eye-finger ray (i.e. aligning eye, finger, and object); while 30 participants (75%) cast a ray orthogonal to the glass surface, as illustrated in Fig. 3.

3.3 Conclusion

75% of the participants that used a raycasting, adopted an orthogonal ray as the most “walk-up and use” technique in a public context. One possible explanation for this is that this approach eliminates the parallax effect since the finger, eye, and object do not need to be aligned. Therefore, in our subsequent work, we consider such interaction and do not try to compensate for the parallax effect. Incidentally, this interaction only requires to track the user’s finger position on the glass, making its implementation quite simple using an infrared touch panel. The main challenge then is to assess whether such form of interaction is efficient in terms of precision. To tack this question, we conduct a controlled pointing study.

4 Pointing Study: Design

The goal of this study is to investigate the impact of four factors on pointing at objects behind a glass-case when using a ray cast perpendicular to the glass. These factors are (1) target size, (2) target distance from the glass, (3) presence of spatial references and (4) users’ head position. Additionally, since interaction in public spaces needs to be easy and quick to use, we also investigated how well users perform the first contact with the touch glass, i.e. before any adjustment, in terms of deviation from the target:

this could provide valuable information for interaction designers about targets shape and size (we discuss its implications in the Discussion section). Finally, since our main goal is to explore the performance of such an interactive technique, we did not compare it to any baseline technique.

4.1 Task and Instructions

The task consisted in pointing one physical spherical target positioned on a 2D cardboard at a specific distance from the glass, without any pointing feedback. Users had to press a keyboard button to start a trial, then touch the glass at the expected position of the target, move the finger for adjusting their pointing if needed, and take back the finger from the glass to end the trial. Pressing a keyboard button with their index finger required to focus on something other than the glass case. This disruption is intended to significantly limit any learning effect. To avoid any impact of the user's finger size and to evaluate a worst case scenario, we used a ray diameter of 0.2 mm, the smallest point size visible in our Unity environment. We discuss the impact of ray sizes in the section Discussion.

We asked participants to select the target at land-on as precisely as possible, and, if they did not reach the target, to move the finger as fast as possible to adjust their pointing. They had to use the index finger of their dominant hand. As mentioned earlier, projection, and therefore dynamic feedback, may not be possible in museum glass cases. Hence, the only interaction feedback used in our study is a white dot to indicate the next target to select (displayed next to the target). The dot became green or red to indicate selection success or failure at land-on. This interaction feedback allowed the user to stay focused on the target throughout the experiment. All these steps are summed up in the Fig. 4 and several trials are presented in the adjunct video.

Of course, in a public context, the pointing would result in displaying some domain related information to the visitor (e.g. in a museum context, the name of the creator, the composition of a rock, etc.): this museographic feedback is generally displayed on a separate display but could also be presented in a more direct manner using spatialized sound or an inside glass projection. Anyway, exploring how to render this museographic information is out of the scope of our work.




Fig. 4. The three steps of a trial: (left) the starting press on the keyboard, (center) the land-on pointing on the glass-case and (right) the adjustment gesture on the glass-case. (Color figure online)

4.2 Target Conditions

We evaluated different target sizes, target distances to the glass, target closeness to a spatial reference and the user's head position with regards to the glass case (and hence to the targets). We defined four target sizes, from 1 to 4 cm in diameter, to cover the spectrum of smallest targets found in a museum. Similar sizes have also been examined in earlier work on targets for touch [37–39], deemed reachable from 10.5 mm to 26 mm.

Our targets were positioned at three different distances to the glass (15, 35 and 55 cm) as a proxy for close, medium and long distances.

For each setting, we positioned 10 spherical targets (cf. Fig. 5) on a 2D cardboard. The targets were aligned on 2 horizontal lines, i.e. 5 targets per line, to explore the impact of a spatial reference on pointing. This condition required to fix the targets in the same locations. The top line was approx. in the middle of the glass case (i.e. without any spatial reference), while the bottom line was close to the bottom of the case (acting as a spatial reference).

Finally, the user could be sitting, i.e. with the head in front of the targets, or standing up, i.e. with the head above the targets. We ensured that all participants positioned their head and body at the same height by using a footstool.

4.3 Participants

We recruited 12 participants aged between 21 and 48 ($M = 29.5$, $SD = 8.26$). 11 participants were recruited at the university and were undergraduate students (4), Ph.D. students (4), research assistants (3) or engineer (1). Ten participants were right-handed and two left-handed.

4.4 Design and Procedure

The experiment followed a $4 \times 3 \times 2 \times 2$ within-participant design with target Size (1, 2, 3 or 4 cm diameter sphere), Distance to the glass (15, 35 and 55 cm), Spatial reference (with or without) and User's head position (in front or above). The User's head position was counterbalanced over participants, and Distance and Size were randomly ordered so that two consecutive trials with the same Distance or Size could not appear. The rest of the trials were ordered randomly. The whole experiment lasted between 70 and 110 min.

We created a specific set of 10 targets (containing the different sizes) for training. The training session consisted in selecting every target on this set for every distance and height (i.e. $10 \times 3 \times 2 = 60$ trials). This training lasted approx. 10 min and represents 7,6% of the whole experiment.

In total, we collected $4 \text{ Sizes} \times 3 \text{ Distances} \times 2 \text{ Spatial reference} \times 2 \text{ User's head position} \times 5 \text{ targets} \times 3 \text{ repetitions} = 720$ trials per participant and 8640 trials in total.

4.5 Apparatus

We built a glass case using four plexiglass sheets of $78.5\text{ cm} \times 44\text{ cm} \times 0.5\text{ cm}$ ($L \times H \times T$). For finger tracking, we used an infrared panel (ZaagTech X series v7, $70\text{ cm} \times 40\text{ cm}$, with 40 touch points, resolution of $32\,768 \times 32\,768$, touch response time of 7 ms–13 ms), USB connected to the computer running the study. The available touch interaction surface was $70\text{ cm} \times 40\text{ cm}$. The panel was attached to one side of the glass case. We also reinforced the plexiglass sheet by adding a metal structure all around it to avoid any deformation or movement of the sheet during the interaction. We used two video projectors (to avoid the possible occlusion caused by the participant's arm) connected through a DisplayPort hub – HDMI, to ensure that the visual interaction feedback was displayed even if the user obstructed one of them. We ensured that there was no light reflection on the glass that could hinder interaction. Between each participant, the glass case was cleaned to remove all finger marks. In public context, anti-fingerprint spray can be used to limit any dirt from touches.

We implemented the whole setup on an HP EliteBook laptop running Microsoft Windows 10. Our experiment software was implemented in C# and Unity (version 2018 1.5). Touch events were sent to the application using the TUIO protocol. We modeled the physical environment in Unity to ensure that the interaction feedback was displayed around the target.




Fig. 5. Experimental setup.

4.6 Collected Data and Data Processing

We logged all touch events from keyboard press until finger land-off (from the glass case). As we were interested in understanding how well users perform the first contact (land-on) with the glass case, we logged for each trial the land-on and land-off finger positions and times (i.e. first and last contact with the glass).

We refined the land-on points analysis with two metrics (cf. Fig. 6) already used in [40, 41]: offset and spread. Both metrics rely on the computation of the centroid, i.e. the average land-on position for each target and condition.

The offset is defined as the Euclidean distance (i.e. absolute value) between the center of the target and the centroid, for each condition. In this paper, we also compute the `offsetX` and `offsetY`, which represents the relative distance in x and y between the centroid and the center of the target (i.e. a signed value).

Spread represents the distribution of the land-on points around the centroid, computed as the diameter of the smallest circle containing 95% of all land-on points.

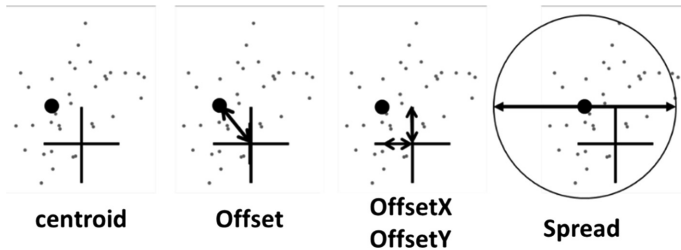


Fig. 6. Illustration of metrics used to analyze land-on points.

Regarding the data analysis, we chose to rely on estimation techniques with 95% Confidence Intervals (CIs) as recommended by the APA [42]. All CIs are 95% BCa bootstrap confidence intervals. Scripts used to compute the geometric average and confidence intervals were used in [43] and are available online [44].

5 Pointing Study Results

We split the analysis into two parts. First, we present the results regarding the target selection task in terms of success rate and time. Then, we analyze the land-on gesture in terms of offset and spread, as explained earlier.

5.1 Target Selection Success Rate

We started by analyzing the results for each target size, independently of the target distance, the user's head position or the spatial reference. Unsurprisingly, the 1 cm targets were the hardest to reach with only 55.8% (CI[48.7%, 65.1%]) of successful trials (cf. Fig. 7). The 2 cm targets were easier to acquire but remain generally difficult, with a success rate of 78.1% (CI[75.8%, 80.6]). The 3 cm and 4 cm targets have a success rate over 90% (3 cm targets: 90.8%, CI[89.5%, 92.4%]; 4 cm targets: 93.2%, CI [90.3%, 95.3%]). For these two larger target sizes, we observed during the experiment that after some successful trials, some users performed the task too fast, i.e. removing the finger before the feedback appeared, leading to some erroneous selections.

Concerning the distance between the glass case and the targets, independently of the other factors, it appears that targets placed 15 cm away were always easier to select (success rate: 82.5% CI[80.1%, 85.4%]) than those placed at 55 cm (success rate: 76.6% CI[73.1%, 80.4%]). However no clear conclusion are revealed when comparing the 15 cm condition to the 35 cm condition (79.3% CI[75.7%, 83%]). This result holds true for all target sizes except for the 1 cm target, which is equally difficult to select for the three distances (i.e. all CIs overlap, cf. Fig. 7).

Concerning the presence of a spatial reference, results strongly establish that targets with a spatial reference (i.e. those on the bottom line) are easier to select (success rate: 83%, CI[80.5%, 86.1%]) than those without spatial reference (success rate: 76%, CI [73.%, 79.8%]). This could be explained by the fact that the spatial reference, i.e. the bottom part of the glass case, facilitates pointing precision. This result is in line with previous work on absolute indirect touch pointing [29].




Fig. 7. Average success rate in % (with 95% CIs) according to target size and distance to the glass.

According to the collected data, the user's head position does not seem to influence the success rate (above: 79.2%, CI[75.9%, 82.7%]; front: 79.8% CI[76.9%, 83.3%]). This result is quite surprising, as we expected the front condition to favor target selection.

5.2 Trial Completion Time

Regarding trial completion time for the different target sizes, results underline the difficulty of selecting the 1 cm targets: in this case, it took users almost 10 s to complete the task (9.4 s, CI[8.1, 10.4]). For the other target sizes (cf. Fig. 8), completion time clearly decreases when the target size increases (2 cm targets: 4.6 s, CI[4.0, 5.1]; 3 cm targets: 3 s, CI[2.6, 3.3]; 4 cm targets: 2.3 s, CI[2.1, 2.6]).

For target distance, we observed the same effect. Although the success rate is very similar for 35 and 55 cm, results strongly establish that on average completion time is smaller when targets are closer to the user (15 cm: 4 s, CI[3.5, 4.6]; 35 cm: 5 s, CI[4.5, 5.4]; 55 cm: 5.33 s, CI[4.6, 6.19], cf. Fig. 8).




Fig. 8. Average time completion in ms (with 95% CIs) according to target size and distance to the glass.

Results did not reveal any effect of spatial reference or user’s head position on completion time.

5.3 Land-on Pointing: Centroid, Offset and Spread

We analyzed land-on pointing results in terms of pointing centroids, offset and spread. In this part, we use 95% confidence interval without bootstrapping.

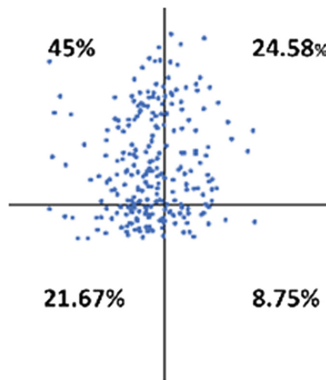


Fig. 9. Distribution of all land-on pointing centroids around the targets. Most of them occur in the top-left area.

Centroid Analysis. Plotting all centroid positions on the same referential gives an overview of pointing at land-on. More than 65% of all centroids are situated on the left of the targets. Almost 70% are situated on the top with 45% on the upper left part (cf. Fig. 9). This result indicates that users had a general tendency to point too high and too far on the left of the physical targets. We further refine this overview by looking at offset, offsetX, and offsetY for our different conditions.

Offset, OffsetX and OffsetY. We found no difference on offset values regarding target size, hence we analysed the results for all target sizes included. Concerning target distance, offset grows with increasing distance (cf. Fig. 10): in comparison to the offset

computed for targets placed at 15 cm (0.66 cm CI[0.55, 0.76])), the offset is 28.8% larger for targets at 55 cm (0.85 cm CI[0.75, 0.95])) and 19.7% larger for targets at 35 cm (0.79 cm CI[0.68, 0.90]) Results do not establish a clear difference between the offset observed for 35 and 55 cm.




Fig. 10. Average offset (left), offsetX (center) and offsetY (right) in cm (with 95% CIs) according to target distance to the glass.

We refine this analysis by looking at offsetX and offsetY. Interestingly, while we can find the same trend for offsetX (i.e. offsetX average is 29.8% larger for targets at 55 cm and 16.4% larger for targets at 35 cm than for those at 15 cm), results reveal no difference in offsetY for targets at different distances (cf. Fig. 10). We can conclude that the distance factor has mostly an impact on land-on precision on the horizontal axis.




Fig. 11. Average offsetX (left) and offsetY (right) in cm (with 95% CIs) according to spatial reference condition.

Regarding the effect of the spatial reference on the offset value, we observed that the average offset is 125% larger for the targets without spatial reference than those with a spatial reference: offset is 1.06 cm (CI[1.02, 1.10]) without spatial reference vs. 0.47 cm (CI[0.44, 0.49]) with a spatial reference. When looking at the impact of target position on offsetX and offsetY, results do not reveal any difference on offsetX. However, offsetY is larger for the targets without spatial reference (0.95 cm CI[0.87, 1.04]) than those with a spatial reference (0.01 cm CI[0.05, 0.04]), as illustrated in Fig. 11.

This analysis can be reinforced when we plot all centroids (cf. Fig. 12) according to both target positions. We can see the large vertical distribution for targets without spatial reference compared to targets with reference.




Fig. 12. Distribution of centroids according to the spatial reference: with (left) or without (right).

Finally, we analyzed the offset according to the User's head position, we found that offset is almost 60% larger when the user's head is in front of the targets (0.65 cm CI [0.61, 0.68]) than above them (0.88 cm CI[0.82, 0.95]). This result holds true for both offsetX and offsetY. This result is surprising since we could expect targets in the front condition to be easier to point, as the user's gaze is aligned with the finger and the target.

Spread. The spread represents the distribution of land-on points around the centroid. Results are in line with our previous findings. Target distance has an impact on spread, which is clearly higher when the distance increases (15 cm: 2.95 cm, CI[2.81, 3.10]; 35 cm: 3.91 cm, CI[3.76, 4.06]; 55 cm: 4.23 cm, CI[4.10, 4.36] cf. Fig. 13).

Furthermore, targets with a spatial reference strongly reduce (cf. Fig. 13) the spread with 3.24 cm (CI[3.15, 3.33]) compared to target without spatial reference (4.16 cm CI [4.02, 4.30]).




Fig. 13. Average spread in cm (with 95% CIs) by target distance (left) and by spatial reference (right).

The target size and user's head position conditions do not seem to have an impact on spread.

5.4 Summary

Our study revealed that both target size and distance to the glass have a strong impact on pointing performance. Very small targets (1 cm) are very hard to point and take a long time to reach. Using larger targets and placing them closer to the glass increases both success rate and completion time. Placing the targets near a spatial reference, such as near the bottom of the glass case, reduces pointing spread and favours pointing performance. A surprising outcome of this study is that the user's point of view has limited impact on pointing performance. We actually observed that when the targets are in front of the user, land-on offset increases.

Overall, our results show that using the glass as touch surface is a promising approach, which allows to easily select targets as small as 3 cm up to 35 cm away from the glass. In the following section, we analyzed the implications of these results for the design of interactive exhibits using touch on a glass.

6 Discussion

Overall, our results show that using the glass as touch surface is a promising approach, which allows to easily select targets as small as 3 cm up to 35 cm away from the glass. In the following section, we analyzed the implications of these results for the design of interactive exhibits using touch on a glass.

6.1 Design Guidelines

Our guidelines can help interaction designers to select the appropriate target size, distance, position and height in the context of absolute indirect pointing on a glass case surrounding a physical object:

- Object size. Designers should favour using orthogonal raycasting from the glass-case for pointing at physical objects placed behind the glass that are larger than 2 cm.
- Object distance to the glass. Objects should be placed at less than 35 cm from the glass. Only objects of 4 cm in size or more, should be placed further away since they still have an average success rate above 90%.
- Object position. Placing objects near the borders of the glass (such as the bottom) provides a spatial reference which improves selection. An approach could be to place smaller objects near a spatial reference, and larger objects which are easier to select in the middle of the case.
- Object and glass case height. Interaction designers can place the objects at any height, from the waist to the head, since our results showed no impact of the user's head position relatively to the object on selection accuracy.

6.2 Improving Target Selection

Our results underline that using a glass as a touch surface to interact with exhibit items is a promising approach, which can be used with relatively small targets placed up to

35 cm away from the glass. Here we present different ways to improve target selection, which could be helpful to facilitate interaction with even smaller and further targets.

Ray Diameter. We chose to study pointing on a physical object by casting a ray with a diameter of 0.2 mm. Since the average offset to the target at land-on ranges from 0.6 to 0.8 cm on average, using a larger diameter will improve target reaching for all conditions. We reran the analysis of our experimental data with two different ray diameters of 1 cm and 2 cm: we found that targets reach on first land-on increases from 42% with a 0.2 mm ray, to 60% with a 1 cm ray and to 73% with a 2 cm ray.

Increase Target Virtual Size. Increasing the virtual size (i.e. the area that can be selected) of targets is a well-known approach [45, 46] to improve pointing performance when the target real size cannot be changed. Our analysis on the spread and the offset shows that land-on points tend to be located above the target. Expanding the virtual size of targets with a vertical rectangular bounding box could facilitate pointing.

6.3 Perspectives

Beyond studying the impact of a different ray diameter or target virtual size, as detailed earlier, a first perspective is to carry a longitudinal in-situ study to analyze the visitors' behavior with such system in a public exhibit. This study would also permit to assess how and where to present the information requested through the direct pointing action on the glass case. As said earlier, the museographic feedback is generally displayed on a separate display but could also be presented in a more direct manner using spatialized sound or an inside glass projection. A second perspective would be to extend our approach to enable selecting volumetric targets. This could be of interest to interact with physical architectural models, which are usually placed behind a glass case. For instance, a model of an airport terminal can be useful for passengers to visualize their way to reach the appropriate terminal. To this end, we need to propose and study solutions for extending our current 2D selection to 3D. One solution will be to adapt existing raycasting disambiguation techniques [16, 47, 48].

7 Conclusion

In this work, we presented a study of *ray casting from a glass surface* for pointing on physical objects behind a glass. This study was motivated by an informative experiment which allowed us to identify the best “walk-up and use” approach for raycasting using an interactive glass case: a raycasting orthogonal to the glass surface. We further studied this particular type of raycasting in a target acquisition experiment. The goal was to evaluate the use of a glass as touch input to point at physical targets placed behind the glass, according to target size, target distance, the presence of spatial references and user's head position with regards to the glass case. Our experimental results reveal that using a glass as touch surface to point at exhibit items behind it is a promising approach, which can be used with relatively small targets placed up to 35 cm away from the glass. Using these results, we proposed several design guidelines for interaction with a physical object behind a glass.

Acknowledgments. This work is partially funded by the French region Occitanie, the neOCampus project (University Toulouse 3) and the AP2 project (ANR grant: AP2 ANR-15-CE23-0001).

References

1. Roberts, J., Banerjee, A., Hong, A., McGee, S., Horn, M., Matcuk, M.: Digital exhibit labels in museums: promoting visitor engagement with cultural artifacts. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, pp. 1–12 (2018). <https://doi.org/10.1145/3173574.3174197>
2. Spindler, M., Dachselt, R.: PaperLens: advanced magic lens interaction above the tabletop. In: Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces, ITS 2009, p. 1 (2009). <https://doi.org/10.1145/1731903.1731948>
3. Martinez Plasencia, D., Berthaut, F., Karnik, A., Subramanian, S.: Through the combining glass. In: Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology, UIST 2014, pp. 341–350 (2014). <https://doi.org/10.1145/2642918.2647351>
4. Schmidt, D., Block, F., Gellersen, H.: A comparison of direct and indirect multi-touch input for large surfaces. In: Gross, T., et al. (eds.) INTERACT 2009. LNCS, vol. 5726, pp. 582–594. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03655-2_65
5. Sears, A., Shneiderman, B.: High precision touchscreens: design strategies and comparisons with a mouse. *Int. J. Man-Mach. Stud.* **34**, 593–613 (1991). [https://doi.org/10.1016/0020-7373\(91\)90037-8](https://doi.org/10.1016/0020-7373(91)90037-8)
6. Forlines, C., Balakrishnan, R.: Evaluating tactile feedback and direct vs. indirect stylus input in pointing and crossing selection tasks. In: Proceedings of ACM CHI 2008 Conference on Human Factors in Computing Systems, vol. 1, pp. 1563–1572 (2008). <https://doi.org/10.1145/1357054.1357299>
7. Argelaguet, F., Andujar, C.: A survey of 3D object selection techniques for virtual environments. *Comput. Graph. (Pergamon)*, **37**, 121–136 (2013). <https://doi.org/10.1016/j.cag.2012.12.003>
8. Mine, M.: *Virtual Environment Interaction Techniques* (1995)
9. Debarba, H.G., Grandi, J.G., Maciel, A., Nedel, L., Boulic, R.: Disambiguation canvas: a precise selection technique for virtual environments. In: Kotzé, P., Marsden, G., Lindgaard, G., Wesson, J., Winckler, M. (eds.) INTERACT 2013. LNCS, vol. 8119, pp. 388–405. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40477-1_24
10. Matulic, F., Vogel, D.: Multiray: multi-finger raycasting for large displays. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, pp. 1–13 (2018). <https://doi.org/10.1145/3173574.3173819>
11. Song, C.G., Kwak, N., Jeong, H.: Developing an efficient technique of Selection and Manipulation in Immersive VE. In: VRST, pp. 142–146 (2000). <https://doi.org/10.1145/502390.502417>
12. Gallo, L., De Pietro, G., Marra, I.: 3D interaction with volumetric medical data: experiencing the Wiimote. In: Proceedings of the 1st international Conference on Ambient Media and Systems, pp. 1–6 (2008). <https://doi.org/10.4108/icst.ambisys2008.2880>
13. Hincapié-Ramos, J.D., Guo, X., Irani, P.: Designing interactive transparent exhibition cases. In: 7th International Workshop on Personalized Access to Cultural Heritage: The Future of Experiencing Cultural Heritage, pp. 16–19 (2014)

14. Teather, R.J., Stuerzlinger, W.: Pointing at 3D targets in a stereo head-tracked virtual environment. In: Proceedings of the IEEE Symposium on 3D User Interfaces, 3DUI 2011, pp. 87–94 (2011). <https://doi.org/10.1109/3dui.2011.5759222>
15. Grossman, T., Wigdor, D., Balakrishnan, R.: Multi-finger gestural interaction with 3D volumetric displays. *ACM Trans. Graph.* **24**, 931 (2005). <https://doi.org/10.1145/1073204.1073287>
16. Grossman, T., Balakrishnan, R.: The design and evaluation of selection techniques for 3D volumetric displays. In: Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology, UIST 2006, p. 3 (2006). <https://doi.org/10.1145/1166253.1166257>
17. Benko, H., Feiner, S.: Balloon selection: a multi-finger technique for accurate low-fatigue 3D selection. In: Proceedings of the IEEE Symposium on 3D User Interfaces, 3DUI 2007, pp. 79–86 (2007). <https://doi.org/10.1109/3dui.2007.340778>
18. Lee, M., Green, R., Billingham, M.: 3D natural hand interaction for AR applications. In: 2008 23rd International Conference Image and Vision Computing New Zealand, IVCNZ (2008). <https://doi.org/10.1109/ivcnz.2008.4762125>
19. Lee, S., Lim, Y., Chun, J.: 3D interaction in Augmented Reality with stereo-vision technique. In: 2013 15th International Conference on Advanced Communication Technology (ICACT), pp. 401–405 (2013)
20. Khamis, M., Buschek, D., Thieron, T., Alt, F., Bulling, A.: EyePACT: eye-based parallax correction on touch-enabled interactive displays. In: Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, vol. 1, pp. 146:1–146:18 (2018). <https://doi.org/10.1145/3161168>
21. Migge, B., Kunz, A.: User model for predictive calibration control on interactive screens. In: Proceedings of the 2010 International Conference on Cyberworlds, CW 2010, pp. 32–37 (2010). <https://doi.org/10.1109/cw.2010.18>
22. Lee, J.H., Bae, S.: Binocular cursor: enabling selection on transparent displays troubled by binocular parallax. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), pp. 3169–3172 (2013). <https://doi.org/10.1145/2470654.2466433>
23. Bandyopadhyay, D., Raskar, R., Fuchs, H.: Dynamic shader lamps: painting on movable objects. In: Proceedings of the IEEE and ACM International Symposium on Augmented Reality, ISAR 2001, pp. 207–216 (2001). <https://doi.org/10.1109/isar.2001.970539>
24. Myers, B.A., et al.: Interacting at a distance: measuring the performance of laser pointers and other devices. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems Changing Our World, Changing Ourselves, CHI 2002, p. 33 (2002). <https://doi.org/10.1145/503376.503383>
25. Freeman, E., Williamson, J., Subramanian, S., Brewster, S.: Point-and-shake: selecting from levitating object displays. In: Proceedings of CHI, pp. 1–10 (2018). <https://doi.org/10.1145/3173574.3173592>
26. Hinckley, K., Wigdor, D.: Input technologies and techniques. In: The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications, pp. 151–168 (2002)
27. Corsten, C., Cherek, C., Karrer, T., Borchers, J.: HaptiCase: back-of-device tactile landmarks for eyes-free absolute indirect touch. In: Proceedings of the ACM CHI 2015 Conference on Human Factors in Computing Systems, vol. 1, pp. 2171–2180 (2015). <https://doi.org/10.1145/2702123.2702277>

28. Gilliot, J., Casiez, G., Roussel, N.: Impact of form factors and input conditions on absolute indirect-touch pointing tasks. In: Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems, CHI 2014, pp. 723–732 (2014). <https://doi.org/10.1145/2556288.2556997>
29. Pietroszek, K., Lank, E.: Clicking blindly: using spatial correspondence to select targets in multi-device environments. In: Proceedings of the 14th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI 2012, pp. 331–334 (2012). <https://doi.org/10.1145/2371574.2371625>
30. Gehring, S., Löchtefeld, M., Daiber, F., Böhmer, M., Krüger, A.: Using intelligent natural user interfaces to support sales conversations. In: Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces, IUI 2012, p. 97 (2012). <https://doi.org/10.1145/2166966.2166985>
31. Green, J., Pridmore, T., Benford, S.: Exploring attractions and exhibits with interactive flashlights. *Pers. Ubiquit. Comput.* **18**, 239–251 (2014). <https://doi.org/10.1007/s00779-013-0661-3>
32. Ridel, B., Reuter, P., Laviolle, J.: The Revealing Flashlight: interactive spatial augmented reality for detail exploration of cultural heritage artifacts. *J. Comput. Cult. Heritage* **7**, 1–18 (2014). <https://doi.org/10.1145/0000000.0000000>
33. Pierce, J.S., Forsberg, A.S., Conway, M.J., Hong, S., Zeleznik, R.C., Mine, M.R.: Image plane interaction techniques in 3D immersive environments. In: Proceedings of the 1997 Symposium on Interactive 3D Graphics, SI3D 1997, p. 39 (1997). <https://doi.org/10.1145/253284.253303>
34. Pierce, J.S., Steams, B.C., Pausch, R.: Voodoo dolls: seamless interaction at multiple scales in virtual environments, pp. 141–145 (1999)
35. Hoang, T.N., Porter, S.R., Thomas, B.H.: Augmenting image plane AR 3D interactions for wearable computers. In: *Conferences in Research and Practice in Information Technology Series*, pp. 9–16 (2009)
36. Marquardt, N., Greenberg, S.: Informing the design of proxemic interactions. *IEEE Pervasive Comput.* **11**, 14–23 (2012). <https://doi.org/10.1109/MPRV.2012.15>
37. Hall, A.D., Cunningham, J.B., Roache, R.P., Cox, J.W.: Factors affecting performance using touch-entry systems: tactual recognition fields and system accuracy. *J. Appl. Psychol.* **4**, 711–720 (1988)
38. Vogel, D., Baudisch, P.: Shift: a technique for operating pen-based interfaces using touch. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2007, p. 657 (2007). <https://doi.org/10.1145/1240624.1240727>
39. Wang, F., Ren, X.: Empirical evaluation for finger input properties in multi-touch interaction. In: Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009, p. 1063 (2009). <https://doi.org/10.1145/1518701.1518864>
40. Holz, C., Baudisch, P.: The generalized perceived input point model and how to double touch accuracy by extracting fingerprints. In: Proceedings of the 28th International Conference on Human Factors in Computing Systems, CHI 2010, p. 581 (2010). <https://doi.org/10.1145/1753326.1753413>
41. Roudaut, A., Pohl, H., Baudisch, P.: Touch input on curved surfaces. In: Proceedings of the International Conference on Human Factors in Computing Systems, CHI 2011, pp. 1011–1020 (2011). <https://doi.org/10.1145/1978942.1979094>
42. VandenBos, G.R.: *Publication Manual of the American Psychological Association*, 6th edn. American Psychological Association, Washington, D.C. (2009)
43. Dubois, E., Serrano, M., Raynal, M.: Rolling-menu: rapid command selection in toolbars using roll gestures with a multi-DoF Mouse. In: Proceedings of the Conference on Human Factors in Computing Systems, April 2018. <https://doi.org/10.1145/3173574.3173941>

44. AVIZ Group: R Macros for Data Analysis. www.aviz.fr/reliefshearing
45. McGuffin, M., Balakrishnan, R.: Acquisition of expanding targets. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Changing our World, Changing Ourselves, CHI 2002, pp. 57–64 (2002). <https://doi.org/10.1145/503376.503388>
46. McGuffin, M.J., Balakrishnan, R.: Fitts' law and expanding targets. *ACM Trans. Comput.-Hum. Interact.* **12**, 388–422 (2005). <https://doi.org/10.1145/1121112.1121115>
47. Kopper, R., Bacim, F., Bowman, D.A.: Rapid and accurate 3D selection by progressive refinement. In: Proceedings of the IEEE Symposium on 3D User Interfaces, 3DUI 2011, pp. 67–74 (2011). <https://doi.org/10.1109/3dui.2011.5759219>
48. Vanacken, L., Grossman, T., Coninx, K.: Exploring the effects of environment density and target visibility on object selection in 3D virtual environments. In: Proceedings of the IEEE Symposium on 3D User Interfaces 2007, 3DUI 2007, pp. 115–122 (2007). <https://doi.org/10.1109/3dui.2007.340783>