

Tiltcasting: 3D Interaction on Large Displays using a Mobile Device

Krzysztof Pietroszek^{1,3}
School of Computing and
Design¹
CSU Monterey Bay
kmpietro@uwaterloo.ca

James R. Wallace^{2,3}
School of Public Health and
Health Systems²
University of Waterloo
james.wallace@uwaterloo.ca

Edward Lank³
Cheriton School of Computer
Science³
University of Waterloo
lank@uwaterloo.ca

ABSTRACT

We develop and formally evaluate a metaphor for smartphone interaction with 3D environments: Tiltcasting. Under the Tiltcasting metaphor, users interact within a rotatable 2D plane that is ‘cast’ from their phone’s interactive display into 3D space. Through an empirical validation, we show that Tiltcasting supports efficient pointing, interaction with occluded objects, disambiguation between nearby objects, and object selection and manipulation in fully addressable 3D space. Our technique out-performs existing target agnostic pointing implementations, and approaches the performance of physical pointing with an off-the-shelf smartphone.

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces

General Terms

Design, 3D Interaction, Large Displays

Author Keywords

3D interaction, mobile interaction, 3D displays, occlusion removal, translation, 3D environments

INTRODUCTION

Recent advances in display technologies have made cost effective, accessible, and mass deployable 3D displays possible. For example, autostereoscopic displays [16] now enable interaction in settings such as airports, building lobbies, and shopping malls. Although interacting with 3D objects in the real world comes naturally to humans, interaction with 3D computer displays continues to be a challenge [2], particularly in settings where users may only casually interact with data such as in public installations or kiosks. For example, humans often have difficulty with tasks such as discovering and selecting objects that are occluded from view [6], distinguishing between nearby objects in dense environments [19], and perceiving an object’s depth [5].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UIST '15, November 08 - 11, 2015, Charlotte, NC, USA
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3779-3/15/11\$15.00
DOI: <http://dx.doi.org/10.1145/2807442.2807471>




Figure 1. Tiltcasting enables 3D interaction with nearby large displays via a Mobile Device.

In supporting interaction in public settings, researchers have investigated various methods of 3D interaction that have inherent limitations. For example, stereoscopic effects rely on eye convergence to convey 3D information, and require a user to stand at a distance from the screen when interacting, excluding the possibility of touch interaction with large displays. Other techniques prioritize ease of use, but require specialized hardware such as a ‘magic wand’ [3] that may prohibit spontaneous use due to configuration constraints. Thus there exists an attendant, unfulfilled need to support interactions that are discoverable and accessible to users.

However, end-users frequently have a device that is convenient and available for interaction: a personal smartphone. This smartphone may serve as an access portal to electronic information and as recent research has explored [13], smartphones can also serve as a convenience device for accessing computation embedded in a user’s environment. Given the appeal of large, public displays as a platform for advertising and infotainment, we wish to explore whether a smartphone can serve as an input device for 3D displays in these settings.

We present the design and validation of Tiltcasting (Figure 1), a metaphor that enables users to interact within a 2D plane that is ‘cast’ from their phone into the 3D space. Our results show that Tiltcasting supports efficient selection regardless of occlusion, nearing established performance benchmarks for physical pointing [8], and enables users to accurately judge target depth. Further, and unlike competitive techniques, Tiltcasting supports both efficient selection and freeform positioning of objects within 3D space.

3D INTERACTION CHALLENGES

There are three primary problems that must be addressed to support precise 3D target acquisition and manipulation: the occluded target problem, the target disambiguation problem, and the depth identification problem.

The Occluded Target Problem

In 3D environments, the occluded target problem arises when, from the perspective of the user, a target is obstructed by another object or objects, inhibiting a user's ability to interact with such objects. Elmqvist and Taigas [6] identify four object interactions that may cause occlusion: proximity, intersection, enclosure, and containment. In cases where a target is partially occluded, most techniques allow for selection, but suffer from speed-accuracy tradeoffs. Common solutions to the occlusion problem are to hide or remove occluding objects, to reposition the viewport, or to distort the interaction space [4]. However, these techniques are typically not integrated with selection mechanisms and few techniques explore the issue of a target being completely occluded from the user's viewpoint for occluded targets.

The Target Disambiguation Problem

The disambiguation problem occurs when interacting with a target amongst nearby distractors. For example, Raycasting techniques usually select the first intersected target. However, oftentimes the desired target may lay behind other targets, making the user's selection ambiguous. A number of Raycasting variants have been proposed to solve the disambiguation problem. For example, Grossman and Balakrishnan [7] and Wyss et al. [19] explore techniques that augment traditional Raycasting with a depth component. While these techniques solve the disambiguation problem, they do so at the expense of slower selection times or more degrees of freedom (DoFs), reducing generalizability.

The Depth Identification Problem

For non-stereoscopic displays, identifying the depth position of a target is often a challenging problem without additional contextual cues, and humans may misjudge the relative depth of two on-screen objects. Techniques to imitate 3D environments on displays include partial occlusion [11], depth of field rendering [12] and linear perspective [5, 18]. Many depth identification techniques have been developed to assist humans with perception of 3D scenes. For example, inferring depth from occlusion [11], or using the cross point of two cast rays [19]. However, inferring depth from occlusion requires scenes to include occluded objects, and crossing rays requires users to manipulate many degrees of freedom.

SMARTPHONES AS PLATFORMS FOR INTERACTION

Having identified challenges for effective 3D interaction, our goal was to explore the capabilities of smartphone-based interaction techniques to provide efficient, ergonomic, and practical methods to overcome those challenges in public settings. We performed an iterative design process that explored alternative models of interaction based on *spatial correspondence* [14] and *virtual pointing* [13].

Spatial correspondence targeting relies on a user's ability to map coordinates between two distinct surfaces. For example, artists, architects, interior designers, and engineers all engage in spatial correspondence targeting when beginning to create a painting, floor plan, or technical drawing where one surface (i.e. a subject, building, or room) is mapped to a corresponding replicate (i.e. a painter's canvas, blueprint, or sketch). To validate the use of spatial correspondence for public display interaction, we conducted an empirical study to explore accuracy of interaction in public spaces [14] and found that spatial correspondence users were able to localize interactions to within 4% of the display area. This high level of accuracy was surprising given the lack of visual feedback on the input device, but suggests that spatial correspondence may be useful in supporting interactions with a large display.

In subsequent work [13], we designed and evaluated a smartphone-based Raycasting technique, called Smartcasting, that uses a phone's pitch and yaw for virtual pointing. Users may interact at different depths along the cast ray via the smartphone's touchscreen enabling the selection of any 3D position within the (x, y, z) coordinate system. Our evaluation of Smartcasting revealed that target selections were comparable to those made with a WiiMote, but suggest two practical limitations. First, while Smartcasting suffers from the occluded target problem, the target disambiguation problem and the depth identification problem. Second, participants experienced fatigue when pointing for extended periods of time; a notable contrast to participants who held the phone with both hands during spatial correspondence trials. With these lessons learned, we develop a novel interaction technique that combines the benefits of spatial correspondence and virtual pointing, called Tiltcasting.

TILTCASTING: PROTOTYPING AND DESIGN

We iterated through a number of prototypes and conducted empirical pilot studies before settling on our final design. In particular, we explored which gyroscopic sensors and magnetic compasses available in modern smartphones should be leveraged in our technique. We pilot tested full 3-axis rotation, 2-axis rotation, or 1-axis rotation with 8 participants, but observed no difference in selection times ($F_{1,15} = 0.549, p = .483, \eta_p^2 = .073$). However, we noticed that users typically reverted to using only one degree of freedom: rotation around the x-axis, or "pitch", and so our final design differs from other smartphone-based 3D interaction techniques [15, 9, 10] and requires control only over rotation about the x-axis (1 DoF, pitch), which when combined with touch input on the device's screen (2 DoF) results in a 3 DoF technique, similar to Balloon Selection [1]. This restriction to 1 DoF appeared to improve the performance of Tiltcasting, and selection times approached those of physical pointing [8].

The Tiltcasting Metaphor

The Tiltcasting metaphor defines a 2D interaction plane inside the 3D control space, dividing the control space into three distinct areas: 1) space behind the interaction plane, where objects are displayed to the user but are not selectable, 2) space intersected by the interaction plane, where objects are both visible and selectable by the user, and 3) space in

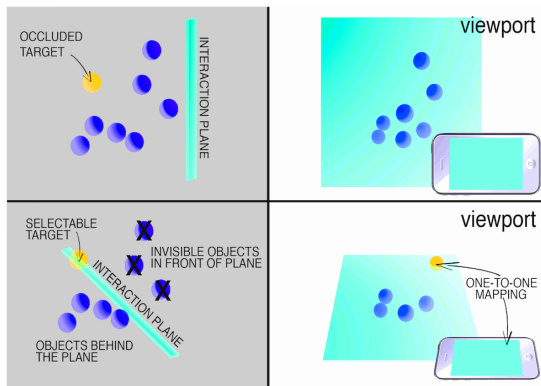


Figure 2. Tiltcasting’s occlusion removal mechanism. a) Yellow sphere occluded by blue sphere, interaction plane is in vertical position. b) Tilting the interaction plane hides three blue objects, revealing the target.

front of the interaction plane, where objects are invisible and not selectable by the user (Figure 2).

Users control the interaction plane within 3D space via their phone’s gyroscope, with rotations about the x-axis corresponding to change in slope for the interaction plane. As the phone rotates, the three defined regions encompass different areas of the 3D space, allowing users to view different parts of the space and to reveal occluded targets. When interacting with a target, the user rotates their phone until the interaction plane intersects the target. A user selects a target by touching the area on the touchscreen which corresponds to the area of contact between the 3D object and the plane. To support selection and dragging, several options present themselves. If targets are sufficiently large, spatial correspondence can support targeting [14]. For smaller targets, in a target-aware implementation the target cut by the plane which is closest to the contact point can be selected, either through a tap or a press and drag. For target agnostic interaction for small targets we support both a tracking and a dragging state as follows: on contact with the surface, a cursor is depicted on the plane and can be repositioned by dragging the finger. To select, a user lifts his or her finger and taps the screen. To drag, a user lifts their finger, touches the screen and drags.

As the interaction plane (and smartphone) changes slope, the visibility of occluded targets may also change. For example, when a user lowers the angle of the phone, objects toward the lower portion of the interaction space may shift from being behind the plane, to intersecting the plane, to above the plane; in turn shifting from being visible, to visible and selectable, to not visible and not selectable. Thus, in selecting an occluded target a user would navigate the interaction plane towards the target, thereby revealing it on-screen by hiding those targets between the interaction plane and the user. Once the position of an occluding object is in front of the viewport it is no longer displayed, and the target is no longer occluded.

Tiltcasting addresses each of the 3D interaction challenges. First, Tiltcasting introduces a novel occlusion removal mechanism: as a user tilts their smartphone, the interaction plane scans through 3D space, removes occluding objects, and reveals potential targets (Figure 2). A particularly useful fea-




Figure 3. Experimental setup

ture is the ability to quickly scan the entire space by swiping the phone down and back up from the vertical position, revealing any occluded targets in the space. Second, Tiltcasting assists with depth identification since the interaction plane provides a depth cue through the linear perspective. Finally, by limiting interaction to a 2D plane within the interaction space, Tiltcasting limits the need to disambiguate nearby targets to only those intersected by the plane. Unlike the other design considerations, our planar interaction space potentially increases the likelihood of the target disambiguation problem occurring compared to Raycasting, where only targets intersected by a line would require disambiguation. We discuss the implications of this choice when interpreting the data collected during our validation.

EMPIRICAL VALIDATION

We compared Tiltcasting to two phone-based Raycasting techniques: target-aware Smartcasting and target-agnostic Smartcasting with Depth Cursor[13]. We identified these Smartcasting implementations as the most appropriate baselines since their use of smartphones reduced potential hardware confounds and had already established performance comparable to common Raycasting implementations such as those on the WiiMote.

Participants

17 participants (11 males, 6 females) participated in the study, whose ages ranged from 19 to 38 ($\bar{x} = 25.2$); 16 right-handed, 1 left-handed. One participant was unable to complete all experimental trials, and their data was excluded from our analysis. Each participant received \$10 remuneration.

Experimental Task

Participants performed a target selection task derived from that used by Vanacken [17]. For each trial, participants first tilted the plane to the vertical position. Then, participants placed the cursor at a start object. Starting with the interaction plane in a vertical position put the Tiltcasting technique at a disadvantage relative to Smartcasting; an optimal initial position of the plane would be 45° , cutting the average selection distance in half. The vertical position of the plane also ensured that the target was fully hidden behind occluder(s) for target-agnostic conditions.

Each trial scene consisted of a start position, destination target, and 45 distractors. The start object was rendered as a yellow sphere, the destination target as a red sphere, and the distractors as blue spheres. Throughout the experiment the start object had a constant size of 1.5 cm, and was displayed in the centre of the 55" display at zero depth along the z-axis. The destination target was placed at a random location on an imaginary sphere with 20 cm radius, with the start target as its centre. Thus, the 3D distance between the start and goal targets was constant across all trials.

Distractors were randomly sized between 1.5 cm and 3.0 cm and randomly positioned such that they did not intersect each other, or the start or goal target. Five distractors were placed around the goal target in a cube-shaped voronoi region [17].

Experimental Design

We used a 2 INTERACTION TECHNIQUE \times 2 TARGET SIZE \times 2 OCCLUSION \times 2 STEREO RENDERING within-subjects design. The study included four independent variables: technique, target size, occlusion and stereo. Targets with either 'small' (0.5°) or 'large' (1.0°) sizes provided two levels of index of difficulty. The experiment environment was rendered with or without stereo. Finally, targets were either fully visible or fully occluded upon starting the trial.

Procedure

Participants were first asked to complete a brief demographic questionnaire. Before the experimental trials, each participant was briefed on each technique, and screened for the ability to see depth. Then the participants completed 5 training trials. These practice trials allowed participants to familiarize themselves with the Tiltcasting technique, and eliminated learning effects from our analysis. Participants then completed 8 blocks of 32 experimental tasks, corresponding to 16 tasks for each target size per block. After the experimental trials, participants completed a post-study questionnaire that examined perceived workload. In total, each session lasted approximately 60 minutes.

Apparatus

We used a 55-inch LG HDTV Cinema 3D stereoscopic display with a pair of passive circularly polarized LG glasses. For input, an iPhone 5 transmitted gyroscope and touch events at 10Hz over a local 802.11n wireless network. Participants were seated 3m in front of a display that was centred to their eye line (Figure 3).

Data Collection and Analysis

All gyroscope and touch interactions were logged to computer files. Selection time was the primary experimental measure, defined as the time taken between entering the start position and reaching the destination target. NASA Task Load Index (TLX) data was collected post-trial. Repeated Measures Analysis of Variance (RM-ANOVA) tests were conducted to examine differences in selection times between target sizes, target visibility, and depth rendering conditions. Friedman tests were used to examine differences in perceived workload measures. An alpha-value of .05 was used for all tests.

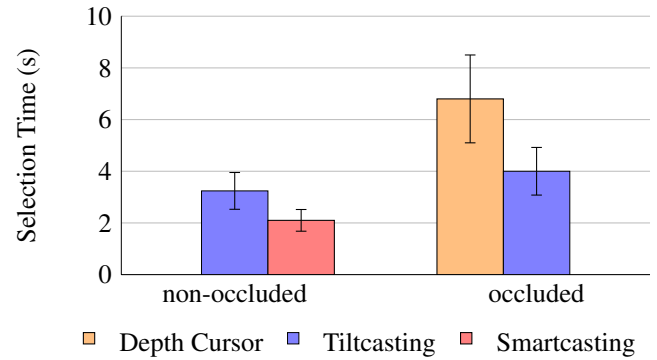


Figure 4. Mean selection times for Tiltcasting, Smartcasting and Smartcasting with Depth Cursor.

RESULTS

We separate our results between target aware and target agnostic techniques. That is, we compare *Smartcasting* and *Tiltcasting*, and *Smartcasting with Depth Cursor* and *Tiltcasting* separately. For each comparison, we evaluated three independent variables: technique, target size, and stereo rendering. Results are summarized in Figure 4 and Table 1.

Non-occluded Targets: Smartcasting vs. Tiltcasting

Participants completed each trial in $2.67s$ ($\sigma = 1.023$) on average, and our analysis revealed a main effect for target selection time ($F_{1,15} = 94.499, p \approx .000, \eta_p^2 = .863$), where Tiltcasting selections took an average of $3.24s$ ($\sigma = .712$) and Smartcasting selections took $2.10s$ ($\sigma = 0.42$) on average. As expected, our analysis revealed main effects for target size ($F_{1,15} = 73.312, p \approx .000, \eta_p^2 = .830$), with small targets taking longer ($3.02s, \sigma = .668$) to select than large targets ($2.33s, \sigma = .416$). Our analysis also revealed a main effect of stereoscopy on selection time ($F_{1,15} = 46.77, p \approx .000, \eta_p^2 = .757$), where selections made with stereoscopic rendering ($3.03s, \sigma = .660$) were slower than those with non-stereoscopic rendering ($2.33s, \sigma = .468$). An interaction effect was found between stereoscopic rendering and size ($F_{1,15} = 6.69, p = .021, \eta_p^2 = .031$), where selection for small targets was faster for non-stereoscopic rendering ($2.57s, \sigma = .528$) than with stereoscopic rendering enabled ($3.47s, \sigma = .872$). For large targets the selection time difference was significant for stereoscopic ($2.58s, \sigma = .532$) vs. non-stereoscopic ($2.08s, \sigma = .424$) rendering. No interaction effect was found between stereoscopic rendering and technique ($F_{1,15} = .063, p = .805, \eta_p^2 = .004$).

Occluded Targets: Smartcasting with Depth Cursor vs Tiltcasting

For occluded targets, participants took $5.40s$ ($\sigma = 2.179$) on average to make selections. Our analysis revealed a main effect for target selection time ($F_{1,15} = 51.781, p \approx .000, \eta_p^2 = .775$), where Depth Cursor selections took an average of $6.8s$ ($\sigma = 1.7$), whereas Tiltcasting selections took an average of $4.0s$ ($\sigma = .922$). Our analysis revealed a main effect of stereoscopy on selection time ($F_{1,15} = 68.62, p \approx .000, \eta_p^2 = .821$), where selections made with stereoscopic rendering ($5.95s, \sigma = 1.18$) were slower than those with

Condition	Target-agnostic				Target-aware			
	Non-Stereoscopic		Stereoscopic		Non-Stereoscopic		Stereoscopic	
	Small	Large	Small	Large	Small	Large	Small	Large
Tiltcasting	3.55s (0.676)	3.18s (0.672)	4.80s (1.18)	4.49s (1.32)	3.15s (.748)	2.62s (.642)	4.05s (1.10)	3.15s (.648)
Smartcasting	—	—	—	—	1.99s (.361)	1.54s (.280)	2.88s (.924)	2.01s (.545)
Depth Cursor	6.48s (2.07)	5.84s (1.48)	8.41s (2.18)	6.10s (1.44)	—	—	—	—

Table 1. Average selection times for Tiltcasting, Smartcasting and Smartcasting with Depth Cursor (standard deviations in parentheses).

non-stereoscopic rendering (4.85s, $\sigma = 1.012$). An interaction effect was found between stereoscopic rendering and size ($F_{1,15} = 7.37, p = .016, \eta_p^2 = .033$), where selection for small targets was faster for non-stereoscopic rendering (5.2s, $\sigma = 1.216$) than with stereoscopic rendering enabled (6.6s, $\sigma = 1.38$). For large targets the selection time difference was significant (4.5s, $\sigma = .952$) for stereoscopic rendering. No interaction effect was found for stereoscopic rendering and technique ($F_{1,15} = 1.284, p = .275, \eta_p^2 = .079$).

Perceived Workload Results

Our analysis did not reveal differences between Tiltcasting and target-aware Smartcasting for mental demand, physical demand, performance, effort or frustration. However a difference in temporal demand was found, where participants reported that Smartcasting was less temporally demanding than Tiltcasting ($p = .004, \chi_{1,15}^2 = 8.33$). Our analysis did reveal differences between target-agnostic techniques, where participants expressed a preference for Tiltcasting in mental demand ($p \approx .000, \chi_{1,15}^2 = 15$), physical demand ($p = .001, \chi_{1,15}^2 = 11.267$), effort ($p = .001, \chi_{1,15}^2 = 11.267$) and for Depth Cursor for temporal demand ($p = .046, \chi_{1,15}^2 = 4$).

DISCUSSION

Our results demonstrate that Tiltcasting effectively supports 3D interaction in public, nearing the 2s selection times typically achieved with physical pointing [8]. Moreover, this interaction is supported via an off-the-shelf smartphone and does not require specialized hardware or camera tracking, reduces fatigue via two-handed input, and enables target selection in target-agnostic settings. For occluded targets, selection times for Tiltcasting were on average 20% faster than those completed with Smartcasting with Depth Cursor. For non-occluded targets Tiltcasting performed only marginally worse than Smartcasting – a technique that benefitted from target-aware selection. Further, for target-agnostic selection Tiltcasting was overwhelmingly preferred by participants and was perceived as requiring less effort to use than Depth Cursor. These results validate its design with respect to the occlusion, depth identification, and target disambiguation problems.

Target Occlusion

Tiltcasting provides effective support for both occluded and non-occluded targets. Tiltcasting provided consistent selection times, averaging 3.6s regardless of whether the target was occluded. Further, two choices made in our experimental design emphasize the importance of these differences. First, in our experimental design target selection times include only the time taken to visually identify and the time taken to select targets. The performance loss between occluded and non-occluded conditions was relatively small, and in practice a 15% increase in selection time may be a worthwhile trade-off when compared to a 3 times increase for Smartcasting. Second, we chose to compare a single Tiltcasting implementation against two Smartcasting implementations: one optimized for target-aware selection (Smartcasting), and one for target-agnostic selection (Depth Cursor). This choice was made to ensure that we held Tiltcasting to a high standard when assessing its performance, but does not reflect compromises that would be made in practice in selecting a single virtual pointing implementation for deployment.

Target Disambiguation

As previously discussed, our choice to allow interaction within a plane potentially exaggerates the target disambiguation problem compared to techniques such as Smartcasting where interaction is restricted to a line. However, our results suggest that Tiltcasting has similar error rates in target selections: errors accounted for less than 0.5% of Tiltcasting trials, compared to 0.4% of Smartcasting trials. We attribute this similarity in error rates to the limited degrees of freedom and occlusion removal properties of the Tiltcasting metaphor, and suggest that these design choices outweighed any increased exposure to the target disambiguation problem while interacting within the interaction plane.

Depth Identification

Participants in our study struggled with stereoscopic rendering, and on average it imposed nearly a 1s penalty on selection times regardless of interaction technique. Selections made with Tiltcasting for non-stereoscopic targets were faster than those performed with stereoscopy ($p = .003$), despite additional depth information being available to participants when using the stereoscopic display. These findings suggest that stereoscopic rendering does not provide an advantage, regardless of which technique was being used. However,

our analysis suggests that Tiltcasting may provide benefits for more difficult cases of the depth identification problem. ‘Large’ targets rendered without stereoscopy represented a worst-case scenario in our study, where the depth identification problem was always present. But, we observed no interaction effect between the target size and stereoscopic rendering for Tiltcasting ($\eta_p^2 = .101$), indicating that the depth confusion effect may be diminished by Tiltcasting’s perspective cue and accounted for less than 10% of the variance in our model. Further, these results were supported by a user preference for Tiltcasting for non-stereo trials.

LIMITATIONS

Our evaluation of the Tiltcasting metaphor demonstrates its utility for 3D pointing, and highlights how its design addresses the three identified challenges for 3D interaction. However, as with any laboratory study, this validation has limitations. For example, we did not evaluate scenarios where many on-screen targets are densely packed within a small angular distance; likely a worst-case scenario for Tiltcasting, as it is with other techniques such as Raycasting. Interaction in these scenarios can be addressed via zooming and/or clutching. Similarly, we conducted our evaluation with a group of University students who may not be representative of the population as a whole. These limitations should be addressed in follow-up research, particularly through an expanded range of experimental tasks and in-the-wild validations.

CONCLUSIONS

We have presented the design and evaluation of a novel 3D interaction metaphor called Tiltcasting that supports interaction with public displays via a user’s smartphone. We validated Tiltcasting for use with both stereoscopic and non-stereoscopic displays, and found that it provides effective support for interaction with both occluded and non-occluded targets. Further, our validation suggests that Tiltcasting provides support for depth identification and effective selections. Our work contributes a deeper understanding of 3D interaction, particularly in the context of the occluded target and depth identification problems, and leverages the ubiquity of mobile, personal devices to reduce barriers to use for affordable, accessible, and commercially available public displays.

ACKNOWLEDGEMENTS

We thank the Natural Sciences and Engineering Research Council of Canada for their support.

REFERENCES

1. Hrvoje Benko and Steven Feiner. 2007. Balloon selection: A multi-finger technique for accurate low-fatigue 3d selection. In *3D User Interfaces, 2007. 3DUI'07. IEEE Symposium on*. IEEE.
2. Kruijff Ernst, LaViola Joseph J., Bowman, Doug A. and Ivan Poupyrev. 2004. *3D user interfaces: theory and practice*. Addison-Wesley.
3. Xiang Cao and Ravin Balakrishnan. 2003. VisionWand: Interaction Techniques for Large Displays Using a Passive Wand Tracked in 3D. In *Proc. UIST*. ACM, 173–182.
4. Zeynep Cipiloglu, Abdullah Bulbul, and Tolga Capin. 2010. A framework for enhancing depth perception in computer graphics. *Proc. APGV 1* (2010), 141.
5. Norman D Cook, Asami Yutsudo, Naoki Fujimoto, and Mayu Murata. 2008. Factors contributing to depth perception: behavioral studies on the reverse perspective illusion. *Spat. Vis.* 21 (2008), 397–405.
6. Niklas Elmqvist and Philippas Tsigas. 2008. A taxonomy of 3D occlusion management for visualization. *IEEE TVCG 14* (2008), 1095–1109.
7. Tovi Grossman and Ravin Balakrishnan. 2006. The design and evaluation of selection techniques for 3D volumetric displays. *Proc. UIST* (2006), 3.
8. Faizan Haque, Mathieu Nancel, and Daniel Vogel. 2015. Myopoint: Pointing and Clicking Using Forearm Mounted Electromyography and Inertial Motion Sensors. In *Proc CHI (CHI '15)*. ACM, 3653–3656.
9. Nicholas Katzakis, Kiyoshi Kiyokawa, and Haruo Takemura. 2013. Plane-Casting: 3D Cursor Control With A Smartphone. In *Proc APCHI*. ACM, 199–200.
10. Nicholas Katzakis, Robert J Teather, Kiyoshi Kiyokawa, and Haruo Takemura. 2015. INSPECT: Extending Plane-Casting for 6-DOF Control. In *3DUI*.
11. Mikko Kytö, Aleksi Mäkinen, Jukka Häkkinen, and Pirkko Oittinen. 2013. Improving relative depth judgments in augmented reality with auxiliary augmentations. *ACM TAP 10*, 1 (Feb. 2013), 1–21.
12. Michael Mauderer, Simone Conte, Miguel Nacenta, and Dhanraj Vishwanath. 2014. Depth perception with gaze-contingent depth of field. *Proc. CHI* (2014), 217–226.
13. Krzysztof Pietroszek, Anastasia Kuzminykh, James Wallace, and Edward Lank. 2014. Smartcasting : A Discount 3D Interaction Technique for Public Displays. In *Proc. OzChi'14*. ACM, 1–10.
14. Krzysztof Pietroszek and Edward Lank. 2012. Clicking Blindly: Using Spatial Correspondence to Select Targets in Multi-device Environments. In *Proc. MobileHCI*. ACM, 331–334.
15. Emanuel Sachs, Andrew Roberts, and David Stoops. 1991. 3-draw: A tool for designing 3d shapes. (1991).
16. A R Travis. 1990. Autostereoscopic 3-D display. *Applied Optics 29* (1990), 4341–4342.
17. Lode Vanacken, Tovi Grossman, and Karin Coninx. 2007. Exploring the Effects of Environment Density and Target Visibility on Object Selection in 3D Virtual Environments. *Proc 3DUI* (2007), 117–124.
18. L.R. Wanger, J.A. Ferwerda, and D.P. Greenberg. 1992. Perceiving spatial relationships in computer-generated images. *CG&A 12* (1992).
19. HP Wyss, Roland Blach, and Matthias Bues. 2006. iSith - Intersection-based spatial interaction for two hands. *Proc 3DUI* (2006), 59–61.