Improving Gestural Interaction With Augmented Cursors

Ashley Dover,
G. Michael Poor
Baylor University
Waco, TX 76706, USA
ashley_therriault@baylor.edu
michael_poor@baylor.edu

Darren Guinness
University of Colorado Boulder
Boulder, CO 80309, USA
darren.guinness@colorado.edu

Alvin Jude
Ericsson Research
Santa Clara, CA, USA
alvinjude@acm.org

ABSTRACT

Gesture-based interaction has become more affordable and ubiquitous as an interaction style in recent years. One issue with gestural pointing is the lack of accuracy with smaller targets. In this paper we propose that the use of augmented cursors – which has been shown to improve small target acquisition with a standard mouse – also improves small target acquisition for gestural pointing. In our study we explored the use of Bubble Lens and Bubble Cursor as a means to improve acquisition of smaller targets, and compared it with interactions without them. Our study showed that both methods significantly improved target selection. As part of our study, we also identified the parameters in configuring Bubble Cursor for optimal results.

Keywords

Gestural Interaction; Augmented Cursor

1. INTRODUCTION

Over the last decade, the availability of affordable commercial off-the-shelf components for gestural interaction (i.e. the Xbox Kinect and Leap Motion) have opened the technology to the general public as an acceptable alternative for computer interaction. However, accuracy can be an issue when using mid-air gestures as cursor control or selection mechanisms [10]. The problem becomes increasingly difficult when the user's target is small or clustered with a number of other potential targets. For example, it has been noted that a target width of under 64px dramatically increases selection difficulty [5]. Given how ubiquitous cursor pointing and selection is in everyday computer usage, this problem is one that has been investigated for mouse interaction [12], with multiple methods showing improvement to the traditional interaction. By incorporating some of the more successful solutions with mid-air gestures, we will investigate whether or not the improvements seen when using the mouse can be applicable to this interaction.

SUI '16, October 15-16, 2016, Tokyo, Japan

DOI: http://dx.doi.org/10.1145/2983310.2985765

In this paper, we describe a study to assess the performance (movement time and error rate) and usability (subjective workload) of gestural interaction enhanced with augmented cursors. Using mid-air gesture based interaction, we compared the standard Point Cursor with the Bubble Cursor, which is considered to be an improvement over the standard point cursor [4], and the Bubble Lens, which has been shown to be an improvement of the Bubble Cursor [9]. Once completed, we found that augmented cursors significantly improved the performance of gestural interaction, in both movement time and error rate.

2. RELATED WORK

Gestural Interaction, which has been studied for over 3 decades, leverages gestures from the body to interact with a computer. Various implementations are typically divided into two different types: (1) those that require the user to wear gloves, devices, or specific markers and (2) touchless gestural interaction. The latter leverages the "Come As You Are" design principle, which states that users should not be required to wear devices or specific markers to interact with a system [11].

Recent devices such as the Xbox Kinect, Leap Motion, and Myo Armband have gained popularity amongst researchers. These devices have demonstrated the potential use of gestural interfaces in medical professions that require sterile environments [11], as an accessibility device for those with impairments [1], and in mixed reality environments with head mounted displays [3]. Given the accuracy of this technology and the sensitive nature of usage, it is imperative to find a better form of selection when using this interaction.

Since graphical user interaction became the industry standard, the mouse was the default selection interaction mechanism. While this interaction has changed little over the years in regards to the default interaction, there have been improvements created for selection [12]. At this time, one of the fastest and well known general pointing techniques is the Bubble Cursor [4], in which a target-aware area cursor with a dynamically resizing activation area allows the cursor to always engulf the nearest target. Granted there have been some techniques that have equaled [2] or slightly improved upon [9] the Bubble Cursor's performance, none have surpassed it for general mouse pointing in all cases.

Though it has been shown to improve user interaction, the Bubble Cursor has some limitations. While it works well when a target is surrounded by empty space, the Bubble Cursor resorts to acting like a point cursor when the target is surrounded by other targets in a tight space [4].

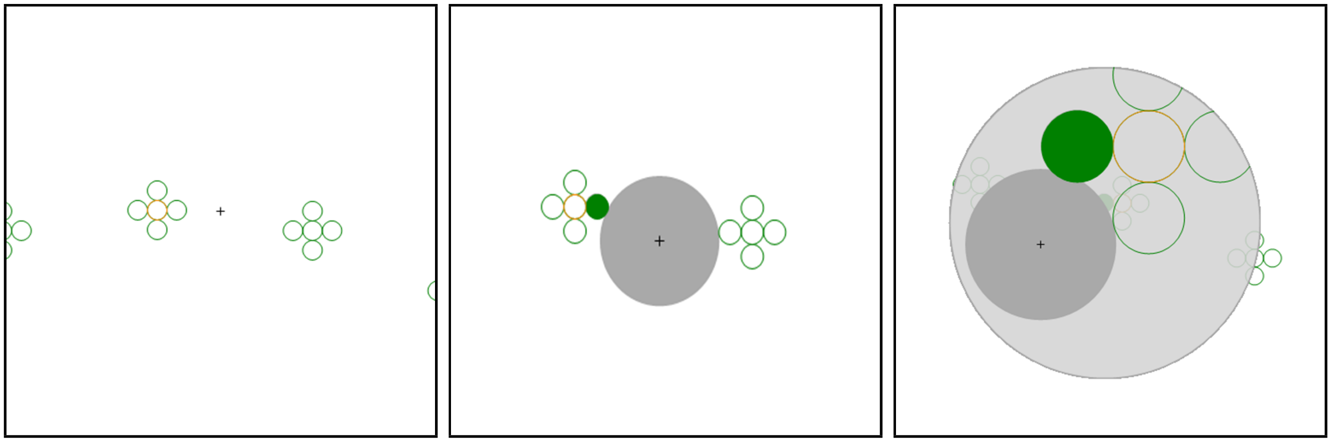

(a) Point Cursor (b) Bubble Cursor (c) Bubble Lens

Figure 1: Cropped screenshots of the task with each of the cursors (Point, Bubble, BubbleLens) showing the intended target in the center (orange outline) surrounded by distractor targets (green).

This type of interaction is very common in today's user interfaces (tool palettes, toolbars, etc). The Bubble Lens [9] was developed with the intention of addressing this limitation; when the target is closely surrounded by other targets, the Bubble Lens will automatically magnify nearby targets using a technique called kinematic triggering. Kinematic triggering is a technique that "continuously examines an unfolding velocity profile to trigger a mode change without explicit invocation [9]." In this work, the authors showed a significant improvement to the Bubble Cursor in both movement time and error rates. A different form of augmented cursors referred to as 'progressive refinement' has also been shown to work with in-air gestural interaction [8].

3. EXPERIMENT

This experiment follows the structure of a previous research [9], except the interaction in our experiment uses mid-air gestures instead of a mouse. The primary focus of this experiment was to assess whether augmented cursors can enhance gestural interaction, and improve performance and usability. Based on our observations of gestural interaction, we hypothesized that the threshold for what is considered a "small" and "dense" target may be different for gestural interaction than from interaction with a mouse. The Bubble Lens cursor uses this threshold to determine when to magnify a target. For mouse interaction this threshold was found to be 10 pixels [9]. We conducted an initial pilot study to determine whether the threshold should be different for gestural pointing, followed by a full study which compared the performance of different cursors with gestural interaction.

3.1 Myo Armband

The Myo Gesture Control Armband was used to recognize the arm's orientation in Euler angles (pitch, yaw, and roll) to control the cursor. In this experiment, we used a modeled interaction using spherical coordinates which has been shown to provide better performance, accuracy, and conformance to Fitts' law [6]. Interactions were performed in a seated position, with participants sitting approximately 3 feet from the monitor. Participants were allowed to perform the mid-air pointing either rested, with the elbow on the desk, or unrested but all participants finished the experiment from a rested position. Before each round, participants performed a calibration stage in which they positioned their arm at the desired angle for each corner of the screen. They were asked to recalibrate the interaction periodically and were allowed to pause and recalibrate if they experienced performance degradation.

3.2 Augmented Cursors

We used 3 different cursor types in our study. The Point Cursor (Figure 1a), used as a control has a single point of activation or hotspot as opposed to area cursors which have larger hotspots which vary in configuration[4]. The Bubble Cursor is a semi-transparent circular area cursor who's size varies to ensure that there is always exactly one target within its hotspot [4]. While the cursor works well when a target is surrounded by empty space, the Bubble Cursor (Figure 1b) resorts to acting like a point cursor when the target is closely surrounded by other targets [4]. The Bubble Lens [9] was developed with the intention of addressing the above limitation; when the target is closely surrounded by other targets, the Bubble Lens (Figure 1c) will automatically magnify nearby targets using a technique called kinematic triggering. Kinematic triggering is a technique that "continuously examines an unfolding velocity profile to trigger a mode change without explicit invocation [9]".

3.3 Experimental Design

Both the pilot and full study focused on comparing the performance of different cursors within gestural interaction used a $3 \times 3 \times 3$ within-subjects design which included 3 cursor types (Point Cursor, Bubble Cursor, Bubble Lens), 3 target widths (diameter of 4, 8, 16 pixels), and 3 spacing levels (zero spacing, half target-width spacing, full target-width spacing). Participants completed the 3 rounds in one of three orderings, based on the Latin cube ordering of interaction type and cursor type (Point, Bubble, Bubble Lens). Each round consisted of 16 trials, the first 3 counting as practice, for each of the 9 effective sizes. Participants completed a total of $3 \times 9 \times 13 = 351$ trials.

	MT (ms)		Error	
	mean	sd	mean	sd
Point Cursor	7570	4486	0.47	0.24
Bubble Cursor	5003	2543	0.32	0.23
Bubble Lens	3301	1084	0.13	0.14

Table 1: Overall means and standard deviation (sd) for movement time (MT) in milliseconds (ms), and error rates. Lower means are better for both.

3.4 Procedure

Each study consisted of one session lasting approximately 60 minutes. All testing was conducted in a lab setting on a 30-inch monitor with 2560 x 1600 resolution. Each study included 3 rounds of gestural interaction using all 3 cursors. Each round began with 3 practice trials, followed by a set of test trials. The task performed by the participants was adapted from that used in the original Bubble Lens study [9]: each trial required the participant to select an orange goal target from a set of gray distractor targets. In order to move on to the next trial, the goal target must be successfully selected. Target selection was done using the keyboard's space bar as the Myo's gestural recognition was found to be less reliable. Participants were told to select the target as quickly as possible while being as accurate as possible. They were also able to ask questions or provide feedback at any point during the experiment. After each round, participants were asked to fill out a Likert scale survey, and at the end of the experiment were asked to rank all rounds performed in order of preference.

3.5 Pilot Study

In order to understand what a small target meant within the context of gestural pointing, similar to previous research [9]. We conducted a pilot as described above with 7 Participants (6 male, 1 female) from a local university. The pilot showed that the Bubble Lens improved performance over the other cursors in both movement time and error rate for effective sizes 16 pixels or below. Thus we allowed the Bubble Lens to deploy for effective sizes 16 pixels or below.

3.6 Full Study

In order to better understand and validate the findings from the pilot we conducted a larger validation study with the same factors as above but with a larger sample size. The study was conducted to determine if the augmented cursor improvements held within a larger statistical sample, and to identify the situations in which the cursors did well.

3.6.1 Participants

Eighteen people participated in the study (12 male, 6 female) and all were university students. The age of participants ranged from 18 to 24, with an average age of 20.5. Fourteen of the participants indicated prior experience with a gestural input device.

3.6.2 Performance Results

Movement time (recorded from first cursor movement until target selection) and error rate (percentage of trials in which the goal target was not successfully selected on the first attempt) were recorded for and are shown in Table 1.

Given that the experiment was within-subjects, a repeated measures analysis of variance for both movement time and error rate were performed. The results for movement time were as followed: significant difference between cursor types (ggF (1.552, 26.383) = 110.933, $p < 0.001$, $\eta_p^2 = 0.867$), a significant difference between effective sizes (ggF (2.230, 37.909) = 76.508, $p < 0.001$, $\eta_p^2 = 0.818$), and a significant difference between effective sizes and cursor type (ggF (3.179, 54.044) = 14.303, $p < 0.001$, $\eta_p^2 = 0.457$). A post-hoc test with Fisher's LSD showed statistical significance between all pairs, with all pairs reporting p<.001. We therefore also measured effect size in Cohen's d with pooled variance to evaluate practical significance. An effect size below .2 is generally considered insignificant, .2 to .5 is considered small, .5 to .8 is considered medium, and d above .8 is considered large. We found an effect size of 0.70 between Point Cursor and Bubble Cursor, an effect size of 1.30 between Point Cursor and Bubble Lens, and an effect size of 0.87 between Bubble Cursor and Bubble Lens.

For error rate there was a significant difference between cursor types (ggF (1.959, 33.296) = 93.270, $p < 0.001$, $\eta_p^2 =$ 0.846), a significant difference between effective sizes (ggF (5.265, 89.512) = 59.708, $p < 0.001$, $\eta_p^2 = 0.778$), and a significant difference between effective sizes and cursor type (ggF (7.274, 123.652) = 4.870, $p < 0.001$, $\eta_p^2 = 0.223$). post-hoc test with Fisher's LSD showed statistical significance between all pairs, with all pairs reporting p<.001. We likewise measured effect size in Cohen's d with pooled variance. This showed an effect size of 0.64 between Point Cursor and Bubble Cursor, an effect size of 1.76 between Point Cursor and Bubble Lens, and an effect size of 1.02 between Bubble Cursor and Bubble Lens.

3.6.3 User Feedback

	Point Cur.		Bubble Cur.		Bubble Lens	
	mean	sd	mean	sd	mean	sd
Mental	8.94	5.20	6.67	4.39	4.61	4.00
Physical	11.00	4.81	7.94	3.93	6.78	3.35
Temporal	8.17	5.66	6.56	4.53	7.11	5.32
Performance	10.17	4.68	6.89	3.51	5.72	3.34
Effort	14.56	4.12	10.94	4.17	8.83	4.46
Frustration	13.89	4.97	7.72	4.20	6.78	5.02

Table 2: Mean and standard deviations of workload ratings (NASA TLX) for each cursor. Lower is better for all.

After each round of the experiment, participants filled out a NASA TLX survey [7]. Workload was assessed using six scales for mental demand, physical demand, temporal demand, performance, effort, and frustration. The means and standard deviation of each feature is shown in table 2. From this table we see that, in general, participants rated Bubble Lens better than the Bubble Cursor, and the Bubble Cursor better than the Point Cursor. Friedman's tests for significance showed the following: mental demand ($\chi^2(2) = 7.62$, $p = 0.022$), physical demand ($\chi^2(2) = 13.0$, $p = 0.002$), temporal demand ($\chi^2(2) = 3.39$, $p = 0.184$), performance ($\chi^2(2) = 15.13$, $p < 0.001$), effort ($\chi^2(2) = 22.12, p < .001$), and frustration ($\chi^2(2) = 21.88, p < .001$). A pairwise test of significance was done with Dunn's test (Table 3). The Bubble Lens was rated significantly better than the Point Cursor in all aspects except temporal demand.

	PC - BC		PC - BL		BC - BL	
	Z	p	Z	p	Z	p
Mental	1.17	.120	2.86	**.002**	1.69	**.046**
Physical	1.84	**.033**	2.83	**.002**	1.00	.161
Temporal	0.78	.219	0.61	.271	-0.17	.435
Performance	2.17	**.015**	3.01	**.001**	0.84	.200
Effort	2.54	**.006**	3.81	**.000**	1.27	.103
Frustration	3.60	**.000**	3.91	**.000**	0.35	.363

Table 3: Pairwise significance testing with Dunn's test for all possible pairs of Point Cursor (PC), Bubble Cursor (BC) and Bubble Lens (BL). Z-statistic and p-values are shown, significance at ($p < .05$) is bolded

4. DISCUSSION

The results of our study showed that enhancing gestural interaction with augmented cursors improves performance, specifically the movement time used to select a target and the error rate. This is similar to previous research in mouse interaction[9]. The results also indicated that the Bubble Lens cursor improved performance more than the Bubble Cursor, as demonstrated with a high effect size. The magnification provided by the Bubble Lens may be a key factor in improving gestural interaction, which tends to be less accurate than mouse interaction, especially for small targets.

Subjective feedback from our participants showed that the augmented cursors did lead to better ratings for gestures, with nearly all workload measures significantly improved. Effort and frustration especially showed large differences between the Point Cursor and the two augmented cursors. The workload ratings show that augmented cursors can improve not only performance for gestural interaction, but also the users' experience, which may aid in making gestures a more practical form of daily computer interaction.

5. CONCLUSION

In this paper, our goal was to assess the performance and usability of gestural interaction that has been enhanced with augmented cursors. We compared the standard Point Cursor, Bubble Cursor, and Bubble Lens cursor. We found that the augmented cursors significantly improved both the movement time and error rate of gestural interaction, with the Bubble Lens cursor showing the best performance.

6. FUTURE WORK

This work has shown the effectiveness of augmented cursors in gestural interaction with small targets, but at what size does the augmented cursor stop being beneficial, if ever? By increasing the size of the targets over the course of experimentation, additional application and conclusion can be made about the potential of this interaction. Additionally the approach works with any device which can sense orientation (3+ DOF), potentially allowing for direct external display control using a smartphone or smartwatch. Future experiments might further investigate the power of magnification in improving gestural interaction. Some individuals, especially those with motor impairments, poor vision, or lower hand-eye coordination, may benefit from higher levels of magnification. The ability to automatically detect and adjust the magnification level may be helpful in such cases. Finally, our experiments were conducted in a desktop setting, but it would make sense to perform a similar form of experiment in AR/VR setting where gestures are more prevalent.

7. REFERENCES

[1] G. Bailly, J. Müller, M. Rohs, D. Wigdor, and S. Kratz. Shoesense: a new perspective on gestural interaction and wearable applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1239–1248. ACM, 2012.

[2] O. Chapuis, J.-B. Labrune, and E. Pietriga. Dynaspot: speed-dependent area cursor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1391–1400. ACM, 2009.

[3] B. Ens, J. D. Hincapié-Ramos, and P. Irani. Ethereal planes: a design framework for 2d information space in 3d mixed reality environments. In *Proceedings of the 2nd ACM symposium on Spatial user interaction*, pages 2–12. ACM, 2014.

[4] T. Grossman and R. Balakrishnan. The bubble cursor: enhancing target acquisition by dynamic resizing of the cursor's activation area. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 281–290. ACM, 2005.

[5] D. Guinness, A. Jude, G. M. Poor, and A. Dover. Models for rested touchless gestural interaction. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, pages 34–43. ACM, 2015.

[6] D. Guinness, A. Seung, A. Dover, G. M. Poor, and A. Jude. Modeling mid-air gestures with spherical coordinates. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*, pages 133–133. ACM, 2015.

[7] S. G. Hart and L. E. Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. In P. A. Hancock and N. Meshkati, editors, *Human Mental Workload*, volume 52 of *Advances in Psychology*, pages 139 – 183. North-Holland, 1988.

[8] R. Kopper, F. Bacim, and D. A. Bowman. Rapid and accurate 3d selection by progressive refinement. In *3D User Interfaces (3DUI), 2011 IEEE Symposium on*, pages 67–74. IEEE, 2011.

[9] M. E. Mott and J. O. Wobbrock. Beating the bubble: using kinematic triggering in the bubble lens for acquiring small, dense targets. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 733–742. ACM, 2014.

[10] D. Vogel and R. Balakrishnan. Distant freehand pointing and clicking on very large, high resolution displays. In *Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 33–42. ACM, 2005.

[11] J. P. Wachs, M. Kölsch, H. Stern, and Y. Edan. Vision-based hand-gesture applications. *Communications of the ACM*, 54(2):60–71, 2011.

[12] J. O. Wobbrock, J. Fogarty, S.-Y. S. Liu, S. Kimuro, and S. Harada. The angle mouse: target-agnostic dynamic gain adjustment based on angular deviation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1401–1410. ACM, 2009.