

Beating the Bubble: Using Kinematic Triggering in the Bubble Lens for Acquiring Small, Dense Targets

Martez E. Mott and Jacob O. Wobbrock

The Information School | DUB Group

University of Washington

Seattle, WA 98195 USA

{memott, wobbrock}@uw.edu

ABSTRACT

We present the *Bubble Lens*, a new target acquisition technique that remedies the limitations of the Bubble Cursor to increase the speed and accuracy of acquiring small, dense targets—precisely those targets for which the Bubble Cursor degenerates to a point cursor. When targets are large and sparse, the Bubble Lens behaves like the Bubble Cursor. But when targets are small and dense, the Bubble Lens automatically magnifies nearby targets, making them larger in both visual- and motor-space. Importantly, magnification is *not* governed by an explicit user-invoked mode-switch. Rather, magnification is activated through *kinematic triggering*, a technique that continuously examines an unfolding velocity profile to automatically trigger mode changes based on observed features. In a first study, we found the Bubble Cursor performed poorly when targets had an effective size smaller than 10 pixels. Using this threshold for the Bubble Lens in a second study, we found that the Bubble Lens significantly outperformed the Bubble Cursor, decreasing movement time by 10.2% and error rates by 37.9%, making the Bubble Lens the fastest current pointing technique.

Author Keywords

Mouse pointing; pointing facilitation; target acquisition; pointing techniques; Bubble Cursor; magnification; zooming; lensing; kinematics.

ACM Classification Keywords

H.5.2. [Information interfaces and presentation]: User Interfaces—*Input devices and strategies*.

INTRODUCTION

Acquiring targets with a mouse cursor remains the dominant form of interaction when operating desktop computer systems. Due to the ubiquity and regularity of mouse pointing in everyday computer use, researchers have developed many pointing facilitation techniques to improve users' pointing performance [7,19,24,28,43]. The current

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 ACM 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.
 CHI 2014, April 26–May 01 2014, Toronto, ON, Canada.
 Copyright 2014 ACM 978-1-4503-2473-1/14/04...\$15.00.
<http://dx.doi.org/10.1145/2556288.2557410>

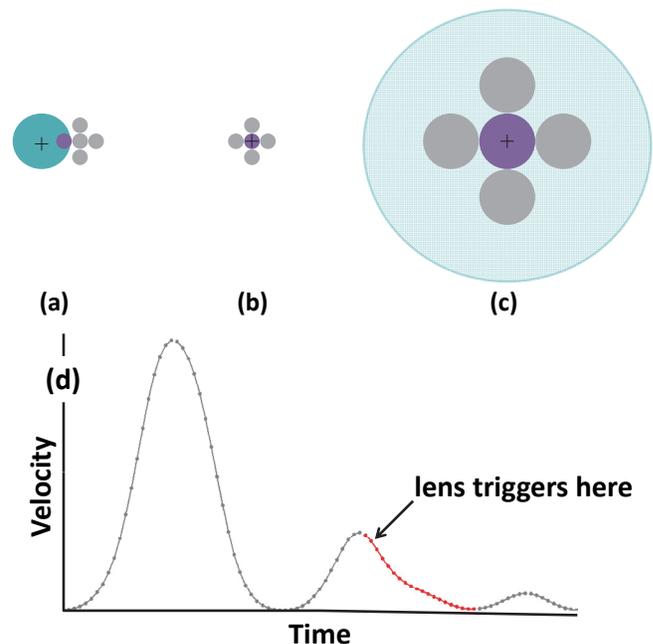


Figure 1. (a) The Bubble Cursor resizes to select the closest target. (b) The Bubble Cursor degenerates to a point cursor in small, dense target situations. (c) Our Bubble Lens automatically magnifies targets, making them larger in visual and motor space. (d) Magnification is triggered on the downward slope of the first corrective submovement in a smoothed velocity profile. The effect is that as a user corrects his motion near the desired target, the lens triggers automatically, making the target easier to acquire.

fastest known general pointing technique is the Bubble Cursor [18], a *target-aware* area cursor with a dynamically resizing activation area that allows the cursor to always engulf the nearest target. Although some techniques have equaled the Bubble Cursor's performance [9], none have surpassed it for general mouse pointing.

Despite its laudable performance, the Bubble Cursor has a key weakness. Although it works well when a target is surrounded by empty space (Figure 1a), it degenerates to a point cursor when a target is densely packed alongside other targets (Figure 1b). The Bubble Cursor's degeneration is problematic, as small, dense targets occur commonly in toolbars, tree views, tool palettes, and other UI designs. In an extreme case where every pixel is a target (*e.g.*, over a paint canvas) [12], the Bubble Cursor degenerates to a single point.

In an effort to strive for better pointing performance, we introduce the *Bubble Lens*, a cursor that remedies the key weakness of the Bubble Cursor and, in so doing, introduces an important reusable technique called *kinematic triggering*. The Bubble Lens generally behaves like the Bubble Cursor, but automatically magnifies nearby targets when the user attempts to acquire small, dense targets (Figure 1c). Unlike most mode-switches, which are explicitly invoked, magnification is activated automatically through *kinematic triggering*, a technique that continuously examines an unfolding velocity profile to trigger a mode change without explicit invocation (Figure 1d). Automatic mode invocation requires very high accuracy or users quickly become frustrated. Our results show that kinematic triggering correctly magnifies 99.3% of the time, sufficient for improving performance and garnering positive user feedback.

In a controlled lab study, we first determined what constitutes “small” and “dense” by discovering at what effective sizes¹ the Bubble Lens improved performance over the Bubble Cursor. We found that the Bubble Cursor performed poorly for targets with an effective size smaller than 10 pixels. In our second study, we found that the Bubble Lens significantly outperformed the Bubble Cursor, decreasing movement time by 10.2% and reducing error rates by 37.9%, making the Bubble Lens the fastest current pointing technique.

The contributions of this paper are: (1) the Bubble Lens invention; (2) the concept of kinematic triggering as a reusable approach to automatic mode-switching based on observed events occurring in a velocity time-series; (3) a novel non-uniform magnification algorithm that, unlike uniform magnification schemes, preserves the distances between target edges while still increasing target sizes; and (4) empirical results showing the Bubble Lens significantly outperforms the Bubble Cursor when acquiring small, dense targets, which are empirically determined to be less than 10 pixels in effective size.

RELATED WORK

A few research projects have incorporated the Bubble Cursor [18] into their designs in hopes of improving target acquisition performance. To improve pen-based pointing accuracy on tablets for older adults, Moffatt and McGrenere created Steadied-Bubbles [35] by combining the Bubble Cursor with Steady Clicks [41]. In their Enhanced Area Cursors for improving pointing performance for users with motor impairments, Findlater *et al.* [15] incorporated the Bubble Cursor into two of their cursor designs, the Motor-Magnifier and the Visual-Motor-Magnifier. These magnifiers were explicitly invoked, and also uniformly magnified underlying space. By contrast, the Bubble Lens

magnifies automatically via kinematic triggering, and magnifies using a custom algorithm that preserves the distance between targets while still enlarging the targets themselves.

In addition to the Bubble Cursor, other target acquisition techniques have incorporated dynamically changing activation areas to improve pointing performance. DynaSpot [9] is a speed-dependent area cursor where the size of the cursor’s activation area is linked to the cursor’s speed. The Cone Cursor and the Lazy Bubble Cursor [30] attempted to strike a balance between the point and the Bubble Cursor by limiting the rate of the Bubble Cursor’s expansion. Both techniques were designed to improve user acceptance over the Bubble Cursor, but neither was able to surpass the Bubble Cursor in terms of pointing performance.

Other techniques investigated by researchers to speed pointing performance include both target prediction and endpoint prediction [4,29,36,40,45]. These related techniques are used to predict either the intended target or the final coordinate of a user’s motion. Unlike these predictive techniques, our kinematic triggering technique does not do any prediction of targets or trajectories. Instead, kinematic triggering looks for an event in a velocity profile to trigger magnification automatically.

Other researchers have investigated the difficulty of acquiring small targets. Chapuis and Dragicevic [8] determined that increasing a small target’s size in motor and visual space will make them easier to acquire. Cockburn and Firth [11] discovered that three target acquisition techniques (expansion [33], sticky [44], and goal-crossing [1,2]) decreased movement time over the default point cursor when acquiring small targets. Cockburn and Brewster [10] showed multimodal feedback (audio, tactile, and pseudo-haptic) could increase the performance of acquiring small isolated targets. Starburst [6] is a technique for improving small, dense target acquisition using a stylus by expanding the size of targets located in target clusters. These prior projects demonstrate the ongoing importance of and interest in improving small-target acquisition.

Past research has shown that lenses can be used to increase target acquisition performance. Ramos *et al.* [39] developed Pointing Lenses, a technique that increases the visual and motor space under a cursor to improve the acquisition of small targets using a stylus. Mankoff *et al.* [32] created a magnification technique to help reduce target ambiguity. The Pointing Magnifier [26] is a pixel-level magnification of contents under an area cursor based on the Visual-Motor-Magnifier created by Findlater *et al.* [15]. Focus + context interfaces [3,17,20,37,38] magnify an area in-place without loss of context, and are an effective means of interacting with objects in multi-scale interfaces.

It is important to note that the Bubble Cursor, and consequently the Bubble Lens, is a *target-aware* [5] pointing technique. Target-aware techniques require the

¹For “effective size,” we will define a value that captures both the size of a target and the target’s density in its immediate vicinity.

cursor to know the locations and dimensions of on-screen targets, and due to the engineering and theoretical challenges they pose [43], few are seen in real-world systems. However, recent innovations in pixel-based reverse engineering of graphical user interfaces [13,14] have brought us closer to bringing these techniques from the lab to the real-world, and the Bubble Cursor has recently been deployed in real-world settings [12]. By extension, the Bubble Lens would be deployable, too; however, engineering such a real-world deployment lies beyond the scope of this initial investigation.

BUBBLE LENS DESIGN

The Bubble Lens is designed to improve the performance of the Bubble Cursor by remedying the Bubble Cursor's performance for small, dense targets. The Bubble Lens achieves this by automatic mode-switching via kinematic triggering, which invokes a lens that shows magnified targets whose distances from one another are preserved. This section describes kinematic triggering, our algorithm preserving targets' edge-to-edge distances under magnification, and how these components help the Bubble Lens achieve superior performance.

Kinematic Triggering for Automatic Lensing

Because they make targets bigger, lenses are known to make targets easier to click [15,39]. Usually, magnifying lenses must be explicitly invoked by the user [15]. Although explicit invocation works fine in situations where pointing speed is not the sole objective (*e.g.*, when magnifying pixels on a paint canvas), explicit invocation takes time and effort. A powerful approach would be for the *automatic* invocation of lenses, but triggering a lens automatically must be done with very high accuracy or users will become frustrated and feel like they are not in control.

When combined with the Bubble Cursor, our concept of *kinematic triggering* uses the unfolding velocity profile of a movement to trigger a lens automatically in the final stages of pointing—early enough to provide a performance benefit, but late enough to ensure that users are trying to enter their intended target. Kinematic triggering is inspired by the *optimized initial impulse model* of Meyer *et al.* [34] (Figure 2). According to this model, a pointing movement consists of a ballistic movement and one or more corrective submovements. The end of a ballistic movement often results in a user undershooting (Figure 2a) or overshooting (Figure 2c) the target. (In rare cases, the end of a ballistic movement can result in a user landing within the bounds of the intended target (Figure 2b).) If a user under- or overshoots, the user must perform corrective submovements (Figures 2d, 2e). Kinematic triggering leverages the optimized initial impulse model to: (1) allow the user to move past distractor targets without triggering a mode change; and (2) to place the lens in a “ready state,” which means the lens is allowed to activate once the cursor has entered a small, dense target cluster. Prior to being

placed in this ready state, the Bubble Lens is ineligible to invoke its lens.

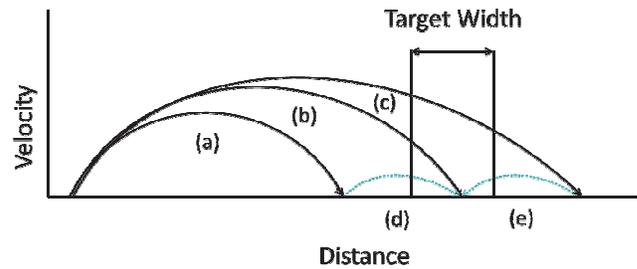


Figure 2. The optimized initial impulse model [34]. Kinematic triggering readies the lens during corrective phases (d) and (e).

To determine when to ready the lens, kinematic triggering continuously examines an unfolding velocity profile. We create a smoothed velocity profile (Figure 3) by temporally resampling the cursor's velocity at 100 Hz, and then smoothing using a Gaussian kernel filter with a standard deviation of 3 [16]. We determine the local minima and maxima in real-time as the smoothed profile unfolds, revealing the location of the ballistic movement and of the corrective submovements.

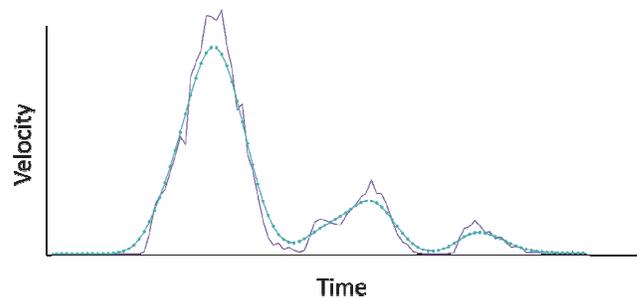


Figure 3. A sample velocity profile. The smoothed profile is shown in blue; the raw profile is shown in purple.

Determining when to put the lens into a ready state is crucial, as activating the lens too early or too late negatively impacts performance. The ideal place along the velocity profile to make the lens eligible for triggering must satisfy two criteria. First, the cursor must travel past the majority of unintended targets without triggering, making any point along the ballistic movement a poor choice for readying the lens. Second, the lens must activate early enough to reduce the amount of time spent in corrective-phase pointing. If it activates too late, the reduction in corrective-phase pointing time will be negligible. Based on these criteria and trial-and-error, we determined that the *start of the downslope of the first corrective submovement* is the best place along the velocity profile to ready the lens—that is, to make it eligible for triggering once the Bubble Lens finds itself in a small, dense target situation (explained next).

Why is the downslope of the first corrective submovement the “right” place? The first corrective submovement occurs after the user has ended the ballistic phase (Figures 2a, 2c), resulting in the user traveling over the majority of

unintended targets, satisfying the first criterion. If a user under- or overshoots during the ballistic phase, they perform a rapid corrective movement toward their target, and begin to decelerate as they near it. This deceleration is the downslope of the first corrective submovement, and results in the user entering the vicinity of their intended target. Readyng the lens during deceleration increases our chance of reducing time spent in corrective-phase pointing, as the magnified targets will be easier to acquire, satisfying the second criterion.

The first corrective submovement is determined by finding the first local maximum that occurs after the peak of the ballistic phase. Because smoothing is done in real-time, the downslope of the first corrective phase cannot be determined instantaneously as it occurs. We found a 30 to 60 millisecond lag occurred before we could reliably determine the cursor had entered the downslope of the first corrective submovement. This lag is small enough to not significantly affect pointing performance, and makes for very accurate triggering of the lens in the vicinity of the target (at 99.3%).

Preserving Targets' Edge-to-Edge Distances when Magnifying

The magnification in the Bubble Lens relies on preserving targets' *edge-to-edge distance*, the intervening space between a pair of targets. In a typical magnification scheme, an entire region is uniformly magnified, including the space between targets (Figures 4a, 4b). In the Bubble Lens, the space between target pairs is preserved, while the targets themselves are nonetheless magnified (Figure 4c). The primary advantage of edge-to-edge distance preservation is it allows magnified targets to remain closer to their original locations in *unmagnified* space, reducing the visual search time for the user once the lens is activated. In essence, edge-to-edge distance preservation avoids an otherwise "jarring" effect for the user.

Edge-to-edge distance preservation is achieved using the following new algorithm. First, all targets are sorted based on their ascending edge-distance from the "trigger target," the target currently engulfed by the cursor at the moment the lens is to be shown. Next, the trigger target is magnified, maintaining the distance between the center of the cursor and the trigger target's nearest edge. The sorted targets are then magnified in closest-to-farthest order from the trigger target.

For each sorted target to magnify, we determine which *previously* magnified target's unmagnified location is closest to the target to magnify. We then magnify this current target using the closest target's magnified location as the reference target. Each target is magnified so that its unmagnified edge-to-edge-distance to this reference target is preserved when magnified. To prevent targets from overlapping, we check for intersecting targets. If targets intersect, the recently magnified target is repositioned using

the intersecting target's magnified location as the new point of reference. Figure 4 shows the effects of edge-to-edge distance preservation.

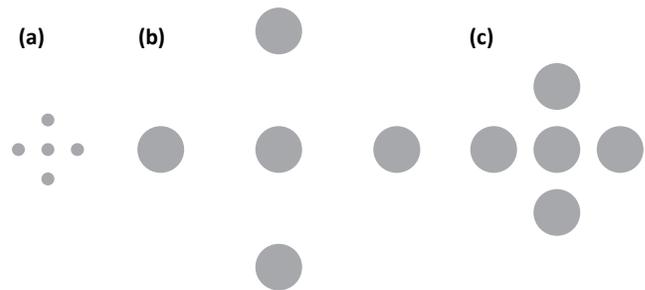


Figure 4. (a) Targets in unmagnified space. (b) Targets under typical uniform 4x magnification. (c) Targets under 4x magnification using our edge-to-edge distance preservation algorithm.

The Bubble Lens: Taking Advantage of Kinematic Triggering and Edge-to-Edge Distance Preservation

The Bubble Lens (Figure 5) builds on the Bubble Cursor [18] to remedy its deficiencies when acquiring small, dense targets. To decrease acquisition time and lower error rates, the Bubble Lens magnifies nearby targets, making them larger in both visual and motor space. Prior work has shown that increasing the size of targets in motor and visual space can help improve target acquisition [8,15].

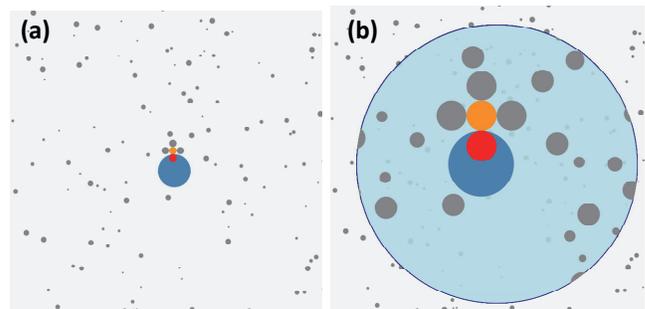


Figure 5. (a) The Bubble Lens before automatic magnification. (b) The Bubble Lens after automatic magnification, which preserves targets' edge-to-edge distances, unlike uniform magnification.

Targets' Effective Size

As previously stated, the Bubble Lens aims to improve pointing performance for small, dense targets. For a performance benefit, the Bubble Lens must magnify a user's intended target *before* the user reaches it. Clearly, the Bubble Lens must be aware of a target's size and density. The Bubble Cursor itself has no notion of target density. (In its determination of the closest target, the Bubble Cursor never actually computes the Voronoi tessellation that geometrically describes each target's effective width [18]. Furthermore, effective width is not an accurate measure of target density anyhow (see Figure 6).)

To determine if targets are dense, we introduce *effective size*. Effective size is our measure of a target's size and a target's density in its immediate vicinity.

$$Effective\ Size = D_A + \frac{\overline{AB} - R_A - R_B}{2} \quad (1)$$

In Equation 1, D_A is the diameter of the currently engulfed target A , \overline{AB} is the distance between the center of the engulfed target A and the center of the target nearest to it B , and R_A and R_B are the radii of the two targets A and B .

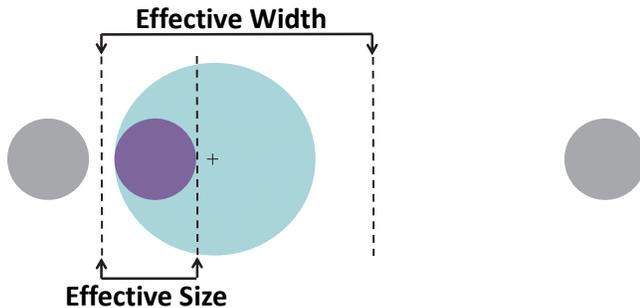


Figure 6. Our effective size calculation is better suited than Voronoi space effective width to capture target size and density. Effective width as defined in [18] is the entire Voronoi region of a target, while our notion of effective size is the size of a target as restricted by its nearest target.

Automatic Triggering of the Lens

A central aspect of the Bubble Lens is that it switches modes automatically. Instead of requiring the user to invoke a mode manually as in prior techniques using magnification [15], the Bubble Lens activates its lens automatically based on kinematic triggering, as described above.

In summary, the lens is activated in three steps. First, kinematic triggering “readies” the lens once the cursor has entered the downslope of the first corrective phase, making eligible the lens to appear. At this point, the cursor should be in the vicinity of the intended target. Second, the edge-distance between the center of the cursor and the currently engulfed target must drop below a minimum threshold. We empirically found that 10 pixels worked well during pilot testing. Third, the effective size of the target according to Equation 1 must be smaller than a threshold, also empirically set at 10 pixels. At the instant these three criteria are met in order, the Bubble Lens triggers its lens automatically, giving the subjective experience of small, dense targets “growing to meet the user” as he approaches.

The purpose of invoking the lens is to reduce small, dense target acquisition time. For targets with an effective size of 10 or more, the lens is not generated, and the Bubble Lens behaves like the default Bubble Cursor. This flexibility allows the Bubble Lens to “kick in” when it is needed while continuing to offer the benefits of the Bubble Cursor when it is not. In essence, the Bubble Lens remedies where the Bubble Cursor is weak, while behaving just like the Bubble Cursor where the Bubble Cursor is strong.

EXPERIMENT

We have described the purpose of the Bubble Lens as remedying the limitations of the Bubble Cursor by improving the acquisition of small, dense targets. But what exactly is “small” and what is “dense?” To find out, we conducted an initial study to determine at what effective

sizes the Bubble Lens improves performance over the Bubble Cursor. After determining the effective size threshold at which the Bubble Lens is helpful, we conducted a validation study to compare the Bubble Lens to the Bubble Cursor.

Participants

Sixteen people participated in the initial study. Seven were female. The average age of participants in the initial study was 27.5 ($SD=2.8$). All but two participants were right-handed. Both left-handed participants said they control a mouse with their right hand in their everyday computer use.

Six people participated in the subsequent validation study. Two were female, and all were right-handed. Their average age was 26.7 ($SD=3.4$).

Apparatus

The experiment testbed was developed in C# .NET 4.5. All sessions were run on a 1920×1080 24” flat panel monitor connected to an iMac 8.1 running Windows 7 at 2.80 GHz with 3.0 GB RAM. The mouse used during the experiment was a Microsoft Comfort Optical Mouse connected over USB. The Windows mouse speed was set at 10, corresponding to an actual C-D gain of about 5 [43], and pointing acceleration was enabled. The testbed captured all mouse cursor movements with millisecond resolution and submillisecond precision.

Procedure

The experiment consisted of a single lab session lasting approximately 35 minutes. At the beginning of each test block, participants completed 7 practice trials to familiarize themselves with the functionality of each cursor. Participants were able to ask any questions during this time. Participants then completed a set of test trials. In each trial, participants were presented with a set of gray distractor targets and one orange goal target (Figure 7). To proceed to the next trial, participants had to successfully select the goal target. All participants were instructed to select the goal target as quickly and accurately as possible—that is, with a balanced approach towards speed and accuracy.

To control target spacing, we placed four distractor targets around the goal target. The immediate distractors were all of equal size and distance from the goal target, the same approach used to evaluate the original Bubble Cursor [18]. The locations and sizes of the remaining distractor targets were randomly generated for each trial, as was the position of the goal target. Six-hundred twenty-five distractor targets were present during each trial. This target amount is halfway between the “sparse” and “clustered” target densities described in prior work [15]. Once all test trials for the cursor type were completed, participants completed a Likert scale questionnaire. At the conclusion of the experiment, participants ranked the cursors in order of preference and provided written feedback.

The Bubble Cursor and Bubble Lens had a maximum radius of 100 pixels. The default magnification factor for the lens was 4, the same as in the Visual-Motor-Magnifier [15]. The generated lens had a size of 500 pixels. In testing, we found this size offered great coverage of the area near the cursor.



Figure 7. A cropped screenshot of the experiment setup. Participants were required to successfully select the orange target in each trial.

Design and Analysis

The initial study was a $3 \times 3 \times 3$ within-subjects design with the following factors and levels:

- *Cursor*: Bubble Lens, Bubble Cursor, point cursor.
- *Width*: Circular diameters of 4, 8, and 16 pixels.
- *Spacing*: no spacing around the target, half target-width spacing, full target-width spacing.

The validation study had the same factors and levels as the initial study, except the point cursor was removed. The levels for our *Width* and *Spacing* factors were used in prior work [15].

Because we wanted to discover whether and where the Bubble Lens increased performance over the Bubble Cursor, we determined the effective size (see Equation 1) for each of the *Width* \times *Spacing* combinations to create an *Effective Size* composite independent variable, resulting in 9 sizes of 4, 5, 6, 8, 10, 12, 16, 20, and 24 pixels.

To counterbalance the presentation of *Cursor*, we used a balanced Latin Square. Nine combinations of *Width* and *Spacing* were blocked by *Cursor*, and were presented in random order. Each block comprised 18 trials, with the first 3 trials serving as practice. Participants in the initial study completed a total of $3 \times 9 \times 15 = 405$ test trials, resulting in a total of 6480 test trials for 16 participants. Participants in the validation study completed a total of $2 \times 9 \times 15 = 270$ test trials, resulting in 1620 test trials for 6 participants.

Temporal measures and submovements were amenable to analyses of variance. Temporal measures were first log-transformed to correct for deviations from normality, and then analyzed with a mixed-effects model analysis of variance [31]. Our model used fixed effects for *Cursor*, *Effective Size*, and *Trial* (nested within *Cursor*); *Subject* was a random effect to accommodate for repeated measures. Although mixed-effects models retain higher denominator degrees of freedom than traditional fixed-effects models, they make detection of significance no easier, and often

harder, due to the use of wider confidence intervals. Mixed-effects models are ideal for input studies involving highly repeated measures correlated across subjects over time.

Other measures required nonparametric analyses. For error rates, we employed the Aligned Rank Transform procedure [22,42]. For Likert responses, we used Friedman and Wilcoxon tests.

RESULTS

This section presents the results of the initial study to determine whether and for what effective target sizes the Bubble Lens improves performance. It also presents results of the validation study that implements this threshold to summatively show a win over the Bubble Cursor.

Initial Study

Movement Time

Movement time was measured from the first mouse move to when the goal target was selected. Mean movement times are shown in Table 1. There was a significant effect of *Cursor* on movement time ($F_{2,6060}=87.91$, $p<.0001$). Not surprisingly, *Effective Size* also exhibited a significant effect on movement time ($F_{8,6060}=457.35$, $p<.0001$), with small effective-size targets taking longer to acquire than large targets. There was a significant *Cursor* \times *Effective Size* interaction ($F_{16,6060}=38.01$, $p<.0001$), because for small, dense targets, the Bubble Lens was fastest, but for large, spacious targets, the Bubble Cursor was fastest. It was exactly these thresholds and crossover points we sought to discover in this initial study. Figure 8 makes them visible.

To formally determine the lower effective-size threshold at which the Bubble Lens offers a benefit over the standard Bubble Cursor, we performed pairwise comparisons. To correct for multiple comparisons, we used Holm's sequential Bonferroni procedure [23]. The pairwise comparisons showed that for targets with effective sizes of 4, 5, 6, and 8, the Bubble Lens was significantly faster than the Bubble Cursor.² As seen in Figure 8, the Bubble Lens was nearly identical to the Bubble Cursor for targets with effective sizes of 10 and 12, and was worse for effective sizes of 16, 20, and 24.

<i>Cursor</i>	<i>MT</i>	<i>Error Rate</i>
<i>Bubble Lens</i>	1546 (382)	2.31% (5.17%)
<i>Bubble Cursor</i>	1577 (437)	3.29% (4.48%)
<i>Point Cursor</i>	1709 (672)	8.61% (9.39%)

Table 1. Overall means for movement time (MT) in milliseconds and error rate (%) in the initial study. Lower is better for both measures. Standard deviations are shown in parentheses.

² In the case of an effective size of 6, the difference was marginal in favor of the Bubble Lens ($F_{1,6060}=3.65$, $p=.056$).

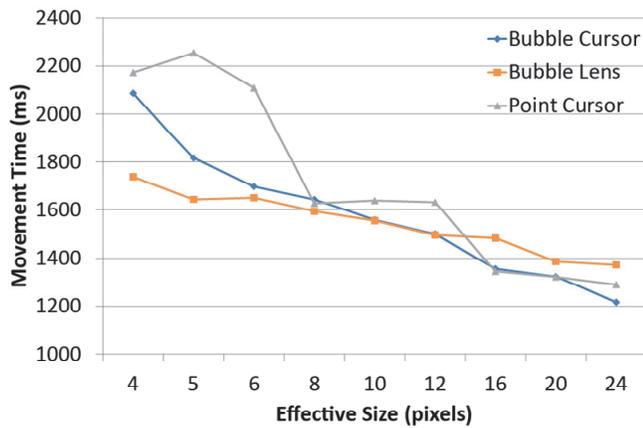


Figure 8. Movement time for each cursor over effective size. The Bubble Lens significantly outperforms the Bubble Cursor for targets with small effective widths (4-8), but performance begins to degrade as effective size grows (16-24).

Error Rate

All participants were required to successfully select the goal target before advancing to the next trial. Therefore, error trials were trials where the first click did *not* result in selecting the goal target. The error rates shown in Table 1 are the percentages of error trials for each cursor.

Mean error rates are displayed in Table 1. There was a significant main effect of *Cursor* on error rate ($F_{2,390}=48.47, p<.0001$). The Bubble Lens was marginally more accurate than the Bubble Cursor ($F_{1,390}=2.73, p=.09$). There was also a significant effect of *Effective Size* on error rate ($F_{8,390}=6.71, p<.0001$), and a significant *Cursor* × *Effective Size* interaction ($F_{16,390}=2.18, p<.01$).

Validation Study

Based on the results found in the initial study, we determined the lens should deploy for targets with effective sizes of up to 10 pixels, and should *not* deploy for larger effective sizes. We added this logic to the Bubble Lens and conducted a validation study to determine whether the Bubble Lens would outperform the Bubble Cursor.

Movement Time

Mean movement times are displayed in Table 2. We found a significant effect of *Cursor* ($F_{1,1345}=211.31, p<.0001$), and *Effective Size* ($F_{8,1345}=148.35, p<.0001$) on movement time. There was also a significant *Cursor* × *Effective Size* interaction ($F_{8,1345}=31.88, p<.0001$). Pairwise comparisons showed that for effective sizes of 4, 5, 6, and 8, the Bubble Lens was significantly faster than the Bubble Cursor ($p<.0001$) (Figure 9). This result confirms the finding from our initial study, which demonstrated that the Bubble Lens was faster for these small, dense targets. The important result here, however, is that despite the Bubble Lens being identical to the Bubble Cursor for effective sizes of 10+, the reduction in movement time for targets with small effective sizes was enough to make the Bubble Lens significantly faster than the Bubble Cursor overall. Put simply, the

Bubble Lens is a faster cursor because it remedies where the Bubble Cursor is weak while retaining where the Bubble Cursor is strong.

<i>Cursor</i>	<i>MT</i>	<i>Error Rate</i>
<i>Bubble Lens</i>	1418 (250)	2.22% (3.43%)
<i>Bubble Cursor</i>	1579 (364)	3.58% (4.97%)

Table 2. Means for movement time (MT) in milliseconds and error rate (%) in the validation study. For both measures, lower is better. Standard deviations are shown in parentheses.

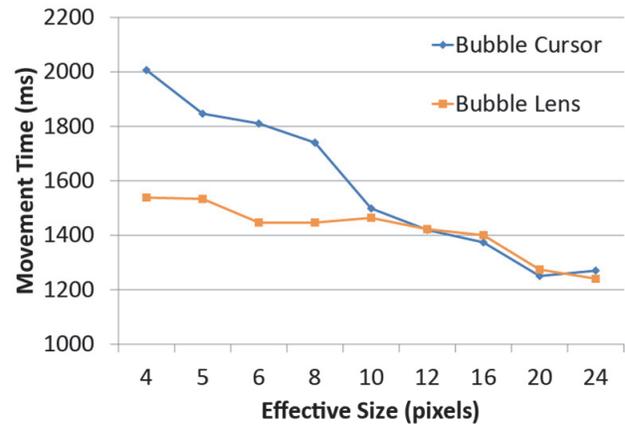


Figure 9. Movement times for the Bubble Lens and the Bubble Cursor. The Bubble Lens outperforms the Bubble Cursor for targets with effective sizes less than 10 pixels. For targets with effective sizes greater than or equal to 10 pixels, the Bubble Lens and Bubble Cursor operated exactly the same.

Error Rate

Error rates for each cursor are displayed in Table 2. Similar to the initial study, the Bubble Lens was marginally more accurate than the Bubble Cursor ($F_{1,85}=3.56, p=.06$). There was neither a significant effect of *Effective Size* ($F_{8,85}=1.15, n.s.$) nor a significant *Cursor* × *Effective Size* interaction ($F_{8,85}=0.81, n.s.$).

The validation study indeed confirmed that the Bubble Lens significantly outperforms the Bubble Cursor in terms of speed, and exhibits a trend towards higher accuracy. Next, we look more closely at our data for when the Bubble Lens and Bubble Cursor behave differently—that is, for when effective sizes are less than 10 pixels—in order to explain the advantages of the Bubble Lens at these sizes.

Submovement Analysis

Submovement analysis offers an insight into pointing behavior. Individual submovements were determined by the number of local maxima in a smoothed velocity profile for each individual trial (see, e.g., Figure 3).

The mean number of submovements were 4.03 ($SD=1.46$) for the Bubble Lens, 4.17 ($SD=2.62$) for the Bubble Cursor, and 4.83 ($SD=3.67$) for the point cursor. These differences were statistically significant ($F_{2,2685}=26.19, p<.0001$). Pairwise comparisons showed that the Bubble Lens did not significantly reduce the number of submovements over the

Bubble Cursor ($F_{1,2685}=1.52$, *n.s.*). However, there was a significant $Cursor \times Effective\ Size$ interaction ($F_{6,2685}=10.60$, $p<.0001$). At the smallest effective size, pairwise comparisons revealed the Bubble Lens significantly reduced the number of submovements over the Bubble and point cursors ($p<.001$).

Bubble Lens Behavior

The lens was triggered for 98.8% of all pointing trials. For targets under the effective size threshold of 10 pixels, the lens was triggered 99.5% of the time. In each trial, we logged whether the goal target appeared in the lens once it was triggered. For 99.3% of trials, the goal target appeared within the lens. This very high percentage shows the high accuracy with which kinematic triggering determined that the lens was eligible for showing.

A crucial aspect of kinematic triggering is preventing the lens from triggering *before* the cursor is near the user's intended target. We were therefore interested in knowing how far away from the goal target did kinematic triggering "ready" the lens. We were also interested in learning how far away from the goal target the lens was actually invoked.

The average straight-line distance from the cursor's starting position to the center of the goal target in each trial was 774 pixels in the Bubble Lens condition. On average, kinematic triggering readied the lens 41.08 pixels away from the goal target, and the lens was invoked 26.54 pixels away from the goal target. Thus, the lens was readied at 94.7% of the distance to the goal target. The lens was invoked at 96.8% of the distance to the goal target.

Subjective Preferences

At the completion of both experiments, each participant was asked to rate the cursors in order from most to least favorite. In the initial study, 9 participants preferred the Bubble Lens, 5 preferred the Bubble Cursor, and 2 preferred the point cursor. In the validation study, 5 of 6 participants preferred the Bubble Lens.

Participants were also asked to provide written feedback describing what they did and did not like about the cursors. Many participants commented on how the Bubble Lens made selecting the small targets easier compared to the Bubble and point cursors. One participant wrote, "I ranked the point cursor last because the smallest targets were so annoying! The Bubble Lens made acquiring smaller targets much easier without much mental effort."

Subjective Workload Ratings

To capture subjective workload, participants filled out a NASA TLX questionnaire [21] after completing all trials for each cursor type (Figure 10). Workload was assessed over six 20-point scales for mental demand, physical demand, temporal demand, performance, effort, and frustration. Tests for the effects of *Cursor* were as follows: mental demand ($\chi^2_{(2,N=16)}=0.06$, *n.s.*); physical demand ($\chi^2_{(2,N=16)}=5.85$, $p=.05$); temporal demand ($\chi^2_{(2,N=16)}=0.29$,

n.s.); performance ($\chi^2_{(2,N=16)}=3.51$, *n.s.*); effort ($\chi^2_{(2,N=16)}=3.96$, *n.s.*); and frustration ($\chi^2_{(2,N=16)}=6.32$, $p<.05$).

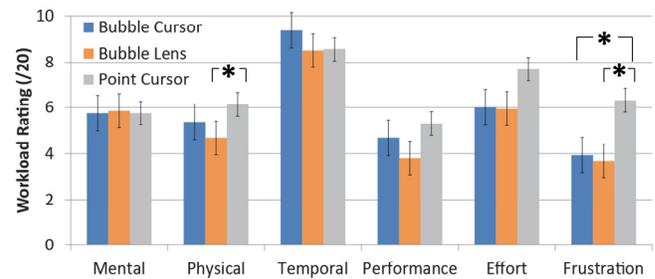


Figure 10. Workload ratings from the NASA TLX. Lower is better for all measures (* $p<.05$).

DISCUSSION

We wanted to discover whether the Bubble Lens could remedy the limitations of the Bubble Cursor by increasing performance when acquiring small, dense targets. First, we had to determine what constitutes a small, dense target for where the Bubble Cursor can be improved upon by the Bubble Lens. Using our definition of *effective size* as a measure of target size and density, we discovered that for targets with an effective size of up to 10 pixels, the Bubble Lens offers significant reductions in movement time and error rate. Consequently, we were able to use an effective size of 10 pixels as a threshold to determine when the lens should trigger. With this threshold employed, the Bubble Lens significantly outperformed the Bubble Cursor in speed, and showed a clear trend towards superior accuracy.

It is interesting to note how, in the initial study, the Bubble Lens negatively affected movement times for targets with effective sizes of 16 and greater. Although we anticipated the Bubble Lens might not perform better than the Bubble Cursor for large effective sizes, we did not anticipate that the Bubble Lens would perform worse than the point cursor also! This result makes sense upon reflection, however, as there is a tradeoff associated with magnifying targets—even automatically doing so. When the magnification occurs, the targets are redrawn, requiring users to visually search for their intended target. Although this search takes only a few moments, we hypothesize that the visual search penalty is greater than the time benefit afforded by having an already-large magnified target. For smaller, denser targets, we believe the benefits outweigh any visual search cost.

Kinematic triggering proved to be an effective method for automatically invoking a mode change, resulting in lenses that showed the desired target 99.3% of the time. (It goes without saying that the lens was *never told* which was the intended target!) It is gratifying that kinematic triggering, by examining only velocity profiles and nothing more, "readied" lenses for deployment about 95% of the way to the intended target—again, with no notion of where the intended target lay. As a reusable concept, kinematic triggering—that is, finding important landmarks in an unfolding smoothed velocity profile and firing or preparing to fire an event—could be implemented in other techniques

that currently require a multi-step target selection process. For example, kinematic triggering could be used to zoom into a document at the end of a scrolling or panning operation, which may result in a technique akin to Speed-Dependent Automatic Zooming [25].

An important aspect of the Bubble Lens's performance is its ability to avoid causing unwanted visual search time on the part of the user by ensuring magnified targets within its lens do not "spread out" as a result of magnification. Keeping targets from moving apart was achieved through our edge-to-edge target distance preservation algorithm. Prior to devising this algorithm, our lens often contained a lot of empty magnified space. We believe our algorithm may be useful in other situations where objects must be magnified but without magnifying the "background" in which they reside. Our algorithm solves the problem of overlap that occurs when retaining magnified objects' original spatial relationships.

Despite the Bubble Lens's 99.3% rate of showing the intended target in the lens, we must still determine how to best de-lens if and when an *unintended* mode change occurs. In the experiment, participants were instructed to right-click to exit an unintended mode change. Although the number of trials where a mode change occurred with the target not in view was tiny (a total of 14 out of 2134 total trials), the average movement time in those trials was 3.9 seconds compared to 1.5 seconds without those trials, a huge performance penalty when the lens gets it wrong.

FUTURE WORK

Future work would include investigating the effectiveness of the Bubble Lens for people with motor impairments. Prior research has shown the Bubble Cursor does not offer significant improvements for people with motor impairments in small, dense target situations [15]. In order to determine whether the Bubble Lens could be effective at reducing acquisition time for users with motor impairments, we would first have to revisit the assumptions made in our kinematic triggering logic. Although it might be true that the downslope of the first corrective submovement works well for people with motor impairments, we cannot know for sure. Also, we would like to test the Bubble Lens with children and older adults. Discovering an automatic way to infer the proper kinematic trigger-point for these various populations would be exciting necessary achievement.

CONCLUSION

We introduced the Bubble Lens, a new target acquisition technique that remedies the limitations of the Bubble Cursor by improving the acquisition of small, dense targets. We discovered that for targets with an effective size up to 10 pixels, the Bubble Lens significantly outperforms the Bubble Cursor. Accordingly, for targets with effective sizes larger than 10 pixels, the Bubble Lens behaves like the default Bubble Cursor. In the process of creating the Bubble Lens, we developed two reusable techniques, (1) kinematic

triggering, and (2) edge-to-edge target distance preservation during magnification, in order to automatically magnify targets near the cursor, increasing their size in both visual and motor space. By pushing our technologies to enable ever-higher levels of human performance, we further not only the limits of our inventions, but also of ourselves.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under grant IIS-0952786. Any opinions, findings, conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect those of any supporter listed above.

REFERENCES

1. Accot, J. and Zhai, S. (1997). Beyond Fitts' law: Models for trajectory-based HCI tasks. *Proc. CHI '97*. New York: ACM Press, 295–302.
2. Accot, J. and Zhai, S. (2002). More than dotting the i's—Foundations for crossing-based interfaces. *Proc. CHI '02*. New York: ACM Press, 73–80.
3. Appert, C., Chapuis, O., and Pietriga, E. (2010). High-precision magnification lenses. *Proc CHI '10*. New York: ACM Press, 273–282.
4. Asano, T., Sharlin, E., Kitamura, Y., Takashima, K., and Kishino, F. (2005). Predictive interaction using the delphian desktop. *Proc UIST '05*. New York: ACM Press, 133–141.
5. Balakrishnan, R. (2004). "Beating" Fitts' law: Virtual enhancements for pointing facilitation. *Int'l J. Human-Computer Studies* 61 (6), 857–874.
6. Baudisch, P., Zotov, A., Cutrell, E., and Hinckley, K. (2008). Starburst: A target expansion algorithm for non-uniform target distributions. *Proc. AVI '08*. New York: ACM Press, 129–137.
7. Blanch, R., Guiard, Y., and Beaudouin-Lafon, M. (2004). Semantic pointing: Improving target acquisition with control-display ratio adaptation. *Proc. CHI '04*. ACM Press, 519–526.
8. Chapuis, O. and Dragicevic, P. (2011). Effects of motor scale, visual scale, and quantization on small target acquisition difficulty. *ACM TOCHI* 18 (3), 1–32.
9. Chapuis, O., Labrune, J.-B., and Pietriga, E. (2009). DynaSpot: Speed-dependent area cursor. *Proc. CHI '09*. New York: ACM Press, 1391–1400.
10. Cockburn, A. and Brewster, S. (2005). Multimodal feedback for the acquisition of small targets. *Ergonomics* 48 (9), 1129–1150.
11. Cockburn, A. and Firth, A. (2003). Improving the acquisition of small targets. *Proc. BCS-HCI '03*. Swindon, England: British Computer Society, 181–196.
12. Dixon, M., Fogarty, J., and Wobbrock, J.O. (2012). A general-purpose target-aware pointing enhancement using pixel-level analysis of graphical interfaces. *Proc. CHI '12*. New York: ACM Press, 3167–3176.
13. Dixon, M. and Fogarty, J. (2010). Prefab: Implementing advanced behaviors using pixel-based reverse

- engineering of interface structure. *Proc. CHI '10*. New York: ACM Press, 1525–1534.
14. Dixon, M., Leventhal, D., and Fogarty, J. (2011). Content and hierarchy in pixel-based methods for reverse engineering interface structure. *Proc. CHI '11*. New York: ACM Press, 969–978.
 15. Findlater, L., Jansen, A., Shinohara, K., Dixon, M., Kamb, P., Rakita, J. and Wobbrock, J.O. (2010). Enhanced area cursors: Reducing fine pointing demands for people with motor impairments. *Proc. UIST '10*. New York: ACM Press, 153–162.
 16. Fisher, R., Perkins, S., Walker, A., and Wolfart, E. (2003). Gaussian smoothing. *Hypermedia Image Processing Reference*. University of Edinburgh.
 17. Furnas, G.W. Generalized fisheye views. (1986). *Proc. CHI '86*. New York: ACM Press, 16–23.
 18. Grossman, T. and Balakrishnan, R. (2005). The Bubble Cursor: Enhancing target acquisition by dynamic resizing of the cursor's activation area. *Proc. CHI '05*. New York: ACM Press, 281–290.
 19. Guiard, Y., Blanch, R., and Beaudouin-Lafon, M. (2004). Object Pointing: A complement to bitmap pointing in GUIs. *Proc. GI '04*. Toronto: CIPS, 9–16.
 20. Gutwin, C. (2002). Improving focus targeting in interactive fisheye views. *Proc. CHI '02*. New York: ACM Press, 267–274.
 21. Hart, S.G. and Staveland, L.E. (1998). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Human Mental Workload*, Hancock and Meshkati (eds). Amsterdam: North Holland, 139–183.
 22. Higgins, J.J. and Tashtoush, S. (1994). An Aligned Rank Transform test for interaction. *Nonlinear World 1*, (2), 2011–211.
 23. Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian J. Statistics 6* (2), 65–70.
 24. Hourcade, J.P., Perry, K.B., and Sharma, A. (2008). PointAssist: Helping four year olds point with ease. *Proc. IDC '08*. New York: ACM Press, 202–209.
 25. Igarashi, T. and Hinckley, K. (2000). Speed-dependent automatic zooming for browsing large documents. *Proc. UIST '00*. New York: ACM Press, 139–148.
 26. Jansen, A., Findlater, L., and Wobbrock, J.O. (2011). From the Lab to the world: Lessons from extending a pointing technique for real-world use. *Proc. CHI EA '11*. New York: ACM Press, 1867–1872.
 27. Kabbash, P. and Buxton, W.A.S. The “Prince” Technique: Fitts' law and selection using area cursors. *Proc. CHI '95*. New York: ACM Press, 273–279.
 28. Kobayashi, M. and Igarashi, T. (2008). Ninja Cursors: Using multiple cursors to assist target acquisition on large screens. *Proc. CHI '08*. New York: ACM Press, 949–958.
 29. Lank, E., Cheng, Y.N., and Ruiz, J. (2007). Endpoint prediction using motion kinematics. *Proc. CHI '07*. New York: ACM Press, 637–646.
 30. Laukkanen, J., Isokoski, P., and Riih a, K.-J. (2008). The Cone and the Lazy Bubble: Two efficient alternatives between the point cursor and the Bubble Cursor. *Proc. CHI '08*. New York: ACM, 309–312.
 31. Littell, R.C., Henery, P.R., and Ammerman, C.B. (1998). Statistical analysis of repeated measures data using SAS procedures. *J. Animal Science 76* (4), 1216–1231.
 32. Mankoff, J., Hudson, S.E., and Abowd, G.D. (2000). Interaction techniques for ambiguity resolution in recognition-based interfaces. *Proc. UIST '00*. New York: ACM Press, 11–20.
 33. McGuffin, M. and Balakrishnan, R. (2002). Acquisition of expanding targets. *Proc. CHI '02*. New York: ACM Press, 57–64.
 34. Meyer, D.E., Smith, J.E.K., Kornblum, S., Abrams, R.A., and Wright, C.E. (1990). Speed-accuracy tradeoffs in aimed movements: Toward a theory of rapid voluntary action. *Attention and Performance 13* (23), 173–226.
 35. Moffatt, K. and McGrenere, J. (2010). Steadied-Bubbles: Combining techniques to address pen-based pointing errors for younger and older adults. *Proc. CHI '10*. New York: ACM Press, 1125–1134.
 36. Pasqual, P.T. and Wobbrock, J.O. (2014). Mouse pointing endpoint prediction using kinematic template matching. *Proc. CHI '2014*. New York: ACM Press. To appear.
 37. Pietriga, E. and Appert, C. (2008). Sigma lenses: focus-context transitions combining space, time and translucence. *Proc. CHI '08*. New York: ACM Press, 1343–1352.
 38. Pietriga, E., Bau, O., and Appert, C. (2009). Representation-independent in-place magnification with sigma lenses. *IEEE TVCG 16* (3), 455–467.
 39. Ramos, G., Cockburn, A., Balakrishnan, R., and Beaudouin-Lafon, M. (2007). Pointing lenses: Facilitating stylus input through visual- and motor-space magnification. *Proc. CHI '07*. New York: ACM Press, 757–766.
 40. Ruiz, J. and Lank, E. (2010). Speeding pointing in tiled widgets: Understanding the effects of target expansion and misprediction. *Proc. IUI '10*. New York: ACM Press, 229–238.
 41. Trewin, S., Keates, S., and Moffatt, K. (2006). Developing Steady Clicks: A method of cursor assistance for people with motor impairments. *Proc. ASSETS '06*. New York: ACM Press, 26–33.
 42. Wobbrock, J.O., Findlater, L., Gergle, D., and Higgins, J.J. (2011). The Aligned Rank Transform for nonparametric factorial analyses using only ANOVA procedures. *Proc. CHI '11*. New York: ACM Press, 143–146.
 43. Wobbrock, J.O., Fogarty, J., Liu, S.-Y. (Sean), Kimuro, S., and Harada, S. (2009). The Angle Mouse: Target-agnostic dynamic gain adjustment based on angular deviation. *Proc. CHI '09*. New York: ACM Press, 1401–1410.
 44. Worden, A., Walker, N., Bharat, K., and Hudson, S. (1997). Making computers easier for older adults to use: Area cursors and sticky icons. *Proc. CHI '97*. New York: ACM Press, 266–271.
 45. Ziebart, B., Dey, A., and Bagnell, J.A. (2012). Probabilistic pointing target prediction via inverse optimal control. *Proc. IUI '12*. New York: ACM Press, 1–10.