Perceptibility and Utility of Sticky Targets

Regan L. Mandryk and Carl Gutwin

University of Saskatchewan

ABSTRACT

Researchers have suggested that dynamically increasing control-to-display (CD) gain can assist in targeting, by increasing the effective width of targets in motor space, which makes targets feel sticky. Although this method has been shown to be effective, there are several unexplored issues that could affect its use in real-world interfaces. One of these is perceptibility: in particular, the difference between the perceptibility and the utility of the technique. If CD gain changes are highly noticeable even at levels that are not helpful, the technique could be seen as overly intrusive. If CD gain changes are more useful than noticeable, however, the technique could be applied more widely. To explore this issue, we carried out a study that tested both the utility and the perceptibility of CD gain in single-target selection tasks. We found that although even small amounts of gain reduction significantly improved targeting times, participants did not consistently notice the effect until the gain difference was much higher. Our results provide new understanding of how changes in CD gain are experienced by users, and provide initial evidence to suggest that sticky targets can benefit users without a high perceptual cost.

KEYWORDS: CD gain, motor space expansion, sticky targets.

INDEX TERMS: H.5.2 Information Interfaces and Presentation: User Interfaces - Interaction Styles; Input devices and strategies

1 INTRODUCTION

Control to display (CD) gain is the relationship between the movement of a physical pointing device such as a mouse and the corresponding movement of a cursor on the screen [8]. With a lower CD gain, the cursor moves more slowly compared to the mouse movement, and with a high gain, the cursor moves more quickly. Most operating systems allow users to set the CD gain, but researchers have suggested that dynamically changing gain, depending on the location of the cursor, can be used to improve selection of targets in graphical interfaces (see Section 2).

In these schemes, the CD gain of the input device is reduced when the cursor is over a target on the screen, slowing the movement of the cursor (Figure 1). This increases the effective width of the target in motor space – that is, a larger movement of the physical device is required to cross the target. Since the size of the target in visual space is unchanged, the effect for the user is that the targets have higher friction, and leads to the term 'sticky targets' for these methods.

Studies have shown that sticky targets can reduce targeting time, but we still know relatively little about how these techniques will affect users. In particular, we do not know how people perceive the change in CD gain – at what level it can be noticed, and how strong the effect feels compared to the utility of the motor-space expansion. This issue is important in considering the technique for use in actual interfaces. If there is a negative gap between utility and perceptibility – that is, if changes to CD gain are highly noticeable at levels that are not helpful, the technique could be seen as overly intrusive and distracting. If, however, CD gain changes are more useful than noticeable, the technique has lower perceptual cost, and can be applied more widely.

To improve understanding of these issues, we carried out a study that explored users' abilities to perceive different levels of CD gain inside targets, and tested their targeting performance at each level. Our study showed that although even small changes to CD gain significantly improved targeting times, participants did not notice the effect until the gain was much higher. These results provide new understanding of how changes in CD gain are experienced, and suggest that sticky targets have considerable potential for providing performance benefits without high costs.

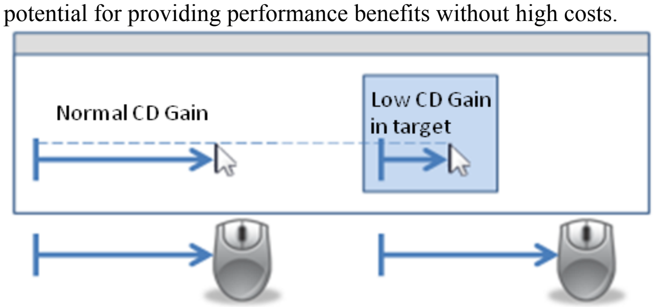

Figure 1. Sticky targets. When the cursor is on the target, twice as much mouse movement is required to cross the same width; therefore the target appears to be 'sticky'. The target expands in motor space, but not in visual space.

2 RELATED WORK

Our investigation is based on three areas of prior work: basic research on targeting, research on target expansion, and research into the perceptibility of changes.

2.1 Targeting Basics

Woodworth [33] first investigated the speed-accuracy tradeoff in aiming movements in 1899, and proposed that aiming movements are comprised of two phases: the initial impulse towards the target, and a deceleration phase under current control to home in on the target. In 1954, Fitts quantified these ideas using a reciprocal tapping task [12], and designed an equation to quantify the difficulty of an aiming task (i.e. the index of difficulty, ID):

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

where A is the amplitude of the movement and W is the width of the target. Fitts also demonstrated that movement time could be correlated with ID and predicted with a regression equation:

$$MT = a + b\ log_2(ID) \quad (2)$$

where a and b are empirically-determined constants. Fitts and Peterson [13] later developed these ideas for discrete aiming movements as Fitts's Law; more recently, Fitts's Law was adapted to model virtual aiming in 2D environments [22], and has been used extensively in HCI to evaluate input devices.

To optimize aiming movements under the constraints of the speed-accuracy tradeoff, most movements consist of a high-velocity ballistic phase (open-loop), and a series of slower, feedback-corrected movements (closed-loop) [24]. MacKenzie et al. [21] examined the velocity profiles of discrete aiming movements and showed that for movements of the same ID (i.e., same movement time), the shape of the velocity profiles differed depending on A and W. The timing and magnitude of the peak velocity relate to the movement amplitude, while the proportion of time spent decelerating depends on the target width.

2.2 Targeting Assistance

To improve targeting time, researchers have created interaction techniques which attempt to 'beat' Fitts's Law in two ways: by reducing the amplitude of movement (A) (e.g. [15]), or by expanding the target (W) (for a review of earlier contributions, see [4]). We consider three types of techniques below, all dealing with target size: techniques that expand the target in visual space, those that expand in both visual and motor space, and those that expand only in motor space (i.e., sticky targets).

2.2.1 Visual-Space-Only Target Expansion

Although we focus primarily on techniques that enlarge the target in motor space, there is also a set of techniques that expand objects only in visual space. Many of these methods use fisheye views, which magnify areas around the cursor without changing the underlying motor size [29]. These visual-only expansions have been shown to be detrimental to targeting performance [16], since they may fool users into thinking that the target is larger than it really is. The experience for users is that targets actually shrink away from the cursor, causing overshooting effects [16].

However, other research has shown that types of visual-only expansion are valuable in targeting. Cockburn and Brock [11] showed that for small targets, visual expansion (that occurs only when the cursor crosses the target's motor boundary) is beneficial because it provides visual feedback that the cursor is on the target.

2.2.2 Motor and Visual Target Expansion

The second type of technique expands targets in both visual and motor space – so that targets both look bigger, and are bigger. These techniques enlarge targets based on the cursor location and are often coupled with endpoint prediction. One of the main issues studied in this area is the amount of time that users need to see the expanded target, as they move towards it, to get the benefit of the expansion. Studies by McGuffin and Balakrishnan [23] and Zhai et al. [35] show that people can make use of the expansion even when it occurs very late in the targeting motion – when the user has already moved 90% of the distance [23], and even when they do not know whether the target will expand or not [35].

There are two main problems with techniques that expand a target's visual space. First, the visual changes made are often highly obvious and thus can be distracting, particularly in cases where the expansion is applied incorrectly. Second, visual expansion must either distort the space [17] or occlude nearby areas of the screen [23, 35]. In situations with sparse targets, this may not be a problem; however, if targets are close together or the underlying data is important, expanding one target makes it more difficult to see and select other objects of interest. Expanding targets in motor space only (described below) does not cause either of these problems.

2.2.3 Motor-Space-Only Target Expansion

Motor-space-only target expansion techniques have generally used disparity between the movement of an input device and the corresponding movement of the on-screen cursor to create the sensation of 'sticky targets'. Although some sticky target approaches have dynamically adapted CD gain to both increase target width and decrease target amplitude, the approach can also be used solely to increase the effective target width. Sticky targets are achieved by slowing the motion of the on-screen cursor relative to the motion of the input device in one of two ways: by adapting CD gain [7, 18, 34]; or by calculating the desired cursor coordinate based on input of the movement device and drawing the cursor at the desired location (cursor warping) [9, 10]. Using both approaches, researchers have found that sticky targets can improve aiming time in both 1D [7, 9, 10] and 2D [18, 34] tasks.

Keyson et al. [18] found movement time improvements for sticky targets compared with normal CD gain, even when the effective ID in movement space was held constant. Worden et al. [34] also implemented sticky icons and found that when this technique was combined with area cursors (i.e. cursors with a larger than normal activation area), targeting performance was improved, especially for older adults. Cockburn and Firth [9] found that the approach was especially helpful for small targets, and Cockburn et al. [10] combined small sticky targets with tactile feedback and non-speech audio to reduce targeting time.

Blanch et al. [7] investigated two levels of target CD gain which resulted in a doubling and quadrupling of target width in motor space. Their approach simultaneously reduced target amplitude, and they showed performance improvements for the sticky targets. In addition, they demonstrated that using ID calculated in movement space to predict movement time modeled aiming more accurately than using the ID of the visual target.

When cursor warping is used both collinear and orthogonal to movement, the resulting sensation is that of a force field, where users 'feel' attractive and repelling forces on their cursors [3, 17]. This force-field approach was used to improve cascading pull-down menu selection times [2], and targeting times in an aiming task [3]. In addition, aiming time to the force-enhanced targets was accurately modeled using a force-adjusted index of difficulty.

Both sticky targets and force fields have been referred to as pseudo-haptic techniques, as they create the illusion of haptic properties by combining the use of a passive input device with disparate visual feedback [19]. Pseudo-haptic approaches have not only been used to assist targeting, but have also: improved users' perception of an object's mass in a virtual environment [19]; represented on-screen textures by creating the sensation of bumps and holes [20]; assisted with boundary scrolling in multi-monitor environments [28]; and supported snapping an object to a target during an alignment task [5].

Most of the previous studies showing targeting benefits of sticky targets investigated single target aiming tasks, while only a handful of studies investigated the effects of neighbouring distracter targets [e.g. 3, 18, 28, 34]. One of the limitations of target-aware CD gain manipulations is that the system must be able to predict the intended target. On a desktop with a few large icons, target prediction may not be difficult, but for realistic applications with clusters of tightly-spaced, small icons, endpoint prediction is not a trivial problem.

2.3 Perceptibility of Change

There is a great deal of prior work on human perception; we do not review this literature here, but we wish to indicate two main themes that inform our study of CD-gain perceptibility.

It is clear that humans can detect very small changes in certain types of visual phenomena, such as jitter in video or animations of moving objects. This suggests that people may be able to easily notice dynamic CD gain manipulation. However, it has also been shown that in other situations, humans fail to notice even large visual differences. Studies of change blindness indicate that people may not notice changes unless they are attending to the visual information being changed. Change blindness has been empirically determined in a variety of experimental conditions, including: a flicker paradigm [27]; changes to an image when 'mudsplashes' appear in non-occluding areas [25]; and gradual changes to a scene without any disruptive influence [30].

Although these studies apply to visual phenomena, analogous results have been shown in the auditory modality (change deafness) [32], and with tactile stimuli [14]. This work suggests that people may not notice changes to CD gain if it is not part of their current attention, or if the changes in CD gain are gradual.

3 STUDY METHODS

CD gain and sticky targets can provide benefits to targeting; however, we do not know much about how people perceive these phenomena, and in order to explore these issues, we carried out a controlled study. Our study was designed to answer four questions about the perceptibility and utility of CD gain on targets:

Q1: Does CD gain assist targeting?
Q2: What levels of gain provide what amount of help?
Q3: At what level is CD gain perceivable?
Q4: Is there a gap between utility and perceptibility?

3.1 Experimental Tasks

Our study used a one-dimensional reciprocal selection task (see Figure 2). This task was chosen to provide stability in the signal, and also because the task's repetitive nature would allow participants more opportunity to notice CD gain changes. Our goal was to choose an experimental task where gain would be as or more obvious to the participant as any targeting task in the real world. This basic task was configured in two ways, to explore two issues: a performance task to determine the utility of CD gain, and a decision task to determine the perceptibility of the effect.

3.1.1 Performance Task - Utility

The goal of the performance task was to determine how different amounts of CD gain affected targeting performance.

The main variable of interest was gain, and we used 11 levels (0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, and 1.0). Level 1.0 corresponds to no gain; that is, the mouse behaved the same way over the target that it did over the other parts of the screen. The other levels indicate the amount of gain that was experienced when the cursor was over the target. For example, gain of 0.5 meant that the cursor moved half as far on the screen for an equivalent movement of the physical mouse. This corresponds to a doubling of the motor-space width of the target. The other gain levels led to different effective target widths: e.g., gain of 0.9 led to an effective width of 1.11 times the normal target width, and gain of 0.05 led to an effective width of 20 times normal width.

We tested three amplitudes (400, 800, and 1600 pixels), and three target widths (50, 100, and 200 pixels) (Figure 2). In this task, participants carried out two blocks of 10 targets per condition; dependent variables were completion time, peak velocity, percent time after peak velocity, errors, and overshoots.

3.1.2 Detection Task - Perceptibility

The second configuration gathered data about whether people could perceive changes in CD gain, and at what level of gain. The task was similar to that described above, but in this task, the CD gain was applied silently to the targets at some point during the ten trials, and the participant's task was to detect the change in CD gain (and press the space bar as soon as they noticed).

The same levels of each factor were used as described above (11 levels of gain, three amplitudes, and three widths). In this task, however, there was only one block of ten trials per condition. The dependent variable was the rate at which people correctly noticed that there was a change in CD gain.

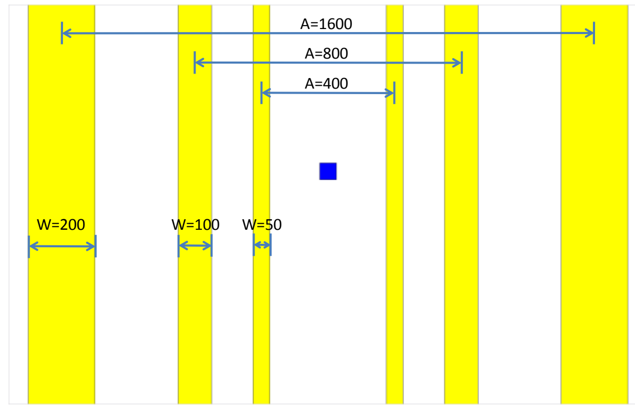

Figure 2. Amplitudes and widths used in the study (all 9 combinations of A and W were used; left and right targets were always the same; there were only 2 targets per trial)

3.1.3 Implementation of CD Gain

Different regions on the screen were set to have different CD gain – the targets had a ratio that was specified by the condition, and the remaining space was always set to a gain of 1.0. The system implemented these different ratios by hiding the system cursor and displaying a custom cursor on the screen. The movement of the custom cursor was calculated by determining how much of each system-cursor movement occurred in each CD gain, and then adjusting the custom cursor accordingly. This approach provides the correct adjustment, regardless of how many mouse events are recorded in the different regions.

3.2 Apparatus

The experiment used a custom study system built in Tcl/Tk (see Figure 2). The system ran on a Windows Vista PC with an Intel Core2 Duo processor at 2.66GHz and 2 GB RAM, a 24-inch monitor set to 1920x1200 resolution, and a gaming-grade Logitech G5 laser mouse, with a resolution of 2000 DPI. Windows pointer acceleration was turned off, and the baseline mouse gain was set to the midpoint, which allowed all tasks to be carried out without clutching.

3.3 Participants and Procedure

Thirteen right-handed subjects (3 females, 8 males) aged 24-34 years participated in the study for a $10 honorarium. Participants all used computers every day, usually for more than 2 hours a day, and all used the mouse as a primary pointing device.

The study took approximately 60 minutes to complete. Participants filled in consent forms and demographic data, and were then introduced to the selection task and CD gain in targets. Participants then performed test trials as described above: the performance task first, then the decision task. In the performance task, participants were instructed to complete the trials as quickly as possible while still maintaining a high degree of accuracy. In the decision task, participants were instructed to press the space bar as soon as they felt a change to the stickiness of the targets.

3.4 Design

The experiment used a three factor (3x3x11) within-subjects full factorial design. The factors were *CD gain* (0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0) (note that the 1.0 level was used three times for a total of 13 levels); *Amplitude* (400, 800, and 1600 pixels); and *Width* (50, 100, and 200 pixels).

The performance task had 117 conditions (13x3x3); with 13 participants and 2 blocks per condition, we collected a total of 3042 data points. The decision task also had 117 conditions but only 1 block per condition, for a total of 1521 data points. In both tasks, levels of CD Gain were presented in a random ordering.

3.5 Data Analyses

To explore the utility of CD gain manipulation, we logged timing information in our application. We also recorded mouse cursor movements, error rates, and the number of times users overshot the target. Mouse cursor x and y positions and the accompanying timestamp were processed trial-by-trial and interpolated to create time-equidistant (120 Hz) arrays. These arrays were then differentiated to provide cursor velocity in both the x- and y-directions. Velocity was calculated according to the formula:

$$Velocity = \sqrt{Vx^2 + Vy^2} \quad (3)$$

The velocity arrays were smoothed using a 25ms moving-average window. From the resulting velocity profiles, we calculated movement time (MT), peak velocity (PkV), and percent of time after peak velocity (%afterPkV) for each trial. Trial 0 of each block was removed for subsequent analyses as the start square was between the two targets, thus trial 0 was half the amplitude of the remaining 9 trials. In addition, we only investigated dependent measures for the second block of trials. In the first block, participants were getting accustomed to the presented combination of CD gain, A, and W. Means and standard deviations across all remaining trials were calculated for each timing-related dependent measure. Outlier trials were defined as trials 3 standard deviations above or below the mean for any given measure; 319 of 11583 trials were removed (2.75%).

A Univariate ANOVA was used to determine if there were any systematic effects of trial and CD gain on MT. Although there were trial effects ($F_{8,96}$=4.8, p<0.001), these effects did not vary by CD gain ($F_{80,979}$=1.2, p>0.05). Our users (all right-handed) were on average faster when aiming from left-to-right than from right-to-left. Because there were no systematic differences of trial and CD gain on MT, we averaged the 9 trials in each block.

Our software logged an error each time a participant clicked outside of the target. Overshoots were logged each time the participant overshot the far vertical boundary of a target. Errors and overshoots were summed across the trials within a block.

We used Fitts's original (1954) method for calculating ID (see equation 1) according to the visual target width in pixels, not the effective width of the motor-expanded target. We conducted a repeated-measures MANOVA with CD gain and ID as factors, MT, PkV, %afterPkV, errors and overshoots as dependent measures, and with α set at 0.05. When the sphericity assumption was violated, the Huynh-Feldt method for adjusting the degrees of freedom was used. All pairwise comparisons used the Bonferroni method for adjusting α to compensate for multiple comparisons.

4 Results

Results for CD gain utility are presented first, followed by results for CD gain perceptibility. CD gain utility is then explored further by examining velocity-based measures and Fitts's Law.

4.1 CD Gain Utility: Movement Time

There was a significant effect of CD gain on MT ($F_{10,120}$=137.9, p<0.001, η^2=0.92). Figure 3 shows how MT decreases as CD gain makes the targets feel stickier. Pairwise comparisons revealed that all levels of CD gain resulted in a significant improvement in MT over no CD gain manipulation (all p<0.01). Although MT decreased linearly with decreases in CD gain, not all neighbouring levels were significantly different from each other. For example, there was no significant MT difference with CD gains of 0.05 or 0.1. For all pairwise comparisons of CD gain levels, see Figure 9. The utility of even small changes to CD gain is shown by the significant differences between no CD gain and CD gain of 0.9 (p<0.01) and 0.8 (p<0.001). These levels represent a target width expansion of only 1.11 and 1.25 times the visible target width, yet result in significant decreases in MT over no CD gain.

As expected, there was also a significant effect of ID on MT ($F_{1.75,21.0}$=750.6, p<0.001, η^2=0.98). Pairwise comparisons showed that each increase in ID significantly increased MT (all p<0.001).

A significant interaction between CD gain and ID on MT ($F_{40,480}$=14.7, p<0.001, η^2=0.55) revealed that the benefits of CD gain on MT are enhanced with increasing ID. Due to the high number of factor levels in the interaction (11 by 5), we do not report all pairwise comparisons, but present the means and standard error of the means in Figure 4.

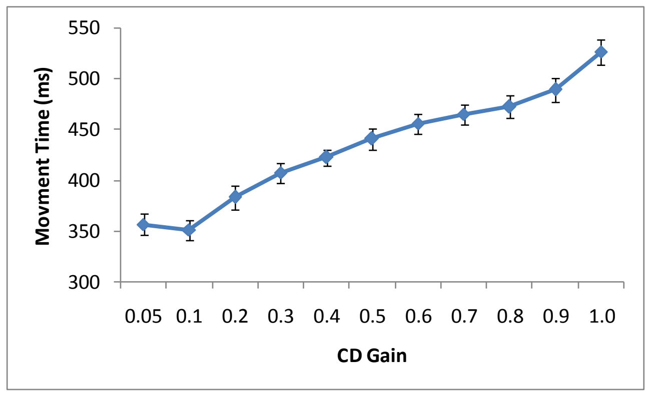

Figure 3. Mean movement times (ms) ± SE by CD gain

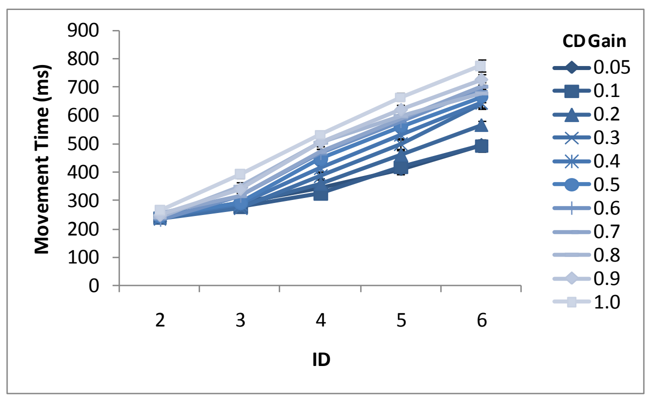

Figure 4. Mean movement times (ms) ± SE by index of difficulty, separated by CD gain.

4.2 CD Gain Utility: Errors and Overshoots

There was a significant effect of CD gain on both errors and overshoots ($F_{5.2, 62.6}$=14.5, p<0.001, η^2=0.55, $F_{6.3,76.0}$=25.0, p<0.001, η^2=0.68 respectively). Both errors and overshoots were significantly decreased with CD gain of 0.7 from CD gain of 1.0. Although both errors and overshoots tended to decrease with CD gain, (see Figure 5), not all pairwise comparisons were significantly different from one another (see Figure 9).

The was also an expected effect of ID (calculated using visual target width) on both errors and overshoots ($F_{4,48}$=15.6, p<0.001,

η^2=0.57; $F_{1.5,17.4}$=64.4, p<0.001, η^2=0.84 respectively). For each increase in ID, the number of overshoots increased significantly (all p<0.02). Although increases in ID also resulted in increased errors, not all pairwise comparisons were significant. The difference between IDs of 2 and 3 was not significant, and there was no difference in the number of errors for IDs of 4, 5 and 6.

These main effects need to be interpreted in light of the significant interaction between CD gain and ID overshoots ($F_{14.5,173.9}$=6.7, p<0.001, η^2=0.36). The effects of CD gain on overshoots were increased for higher IDs (Figure 6). Pairwise comparisons showed no difference in the number of overshoots at different CD values for IDs of 2 or 3, while differences appeared for IDs of 4, 5, and 6. Due to the high number of factor levels, all pairwise comparisons are not presented (see Figure 9); however, means and standard errors are presented in Figure 6.

The interaction between CD gain and ID on the number of errors was not significant ($F_{16.6,199.2}$=1.7, p>0.05, η^2=0.12).

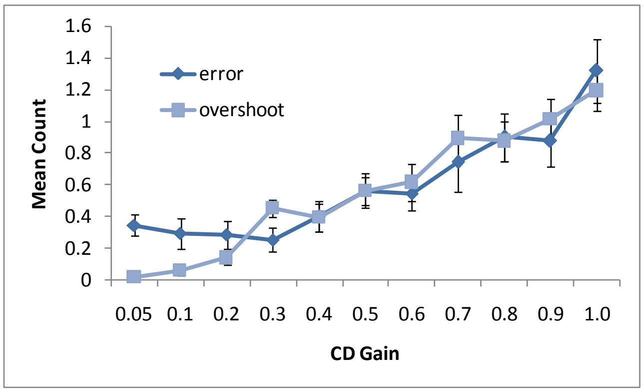

Figure 5. Mean error and overshoots ± SE by CD gain.

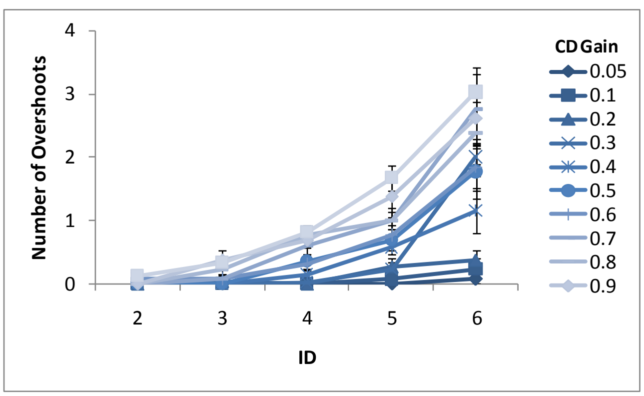

Figure 6. Mean number of overshoots ± SE by index of difficulty, separated by CD gain

4.3 CD Gain Perceptibility: Detection Rate

Our second task tested each participant's ability to detect a change in CD gain that occurred during the block of targeting trials. To sensitize participants to the detection task, we asked them during the first task to rate whether they noticed any CD gain manipulation after each condition (on a scale of 1 to 5). (Note that these results are not a reliable detection measure since participants were biased towards conditions where they performed well).

Our detection task therefore asked participants to press the space bar when target CD gain changed within a block of trials. Due to the detection question in the first experiment, participants were primed to recognize CD gain manipulations. In 23% of the blocks we did not vary CD gain (3 of the 13 levels of CD Gain), and participants were aware that CD gain did not always change.

To determine CD gain detection rates, we first removed all trials where users detected the presence of CD gain manipulation before it was presented. In these false-positive trials, we do not have an indication of whether users realized their mistake if and when CD gain was actually presented, or whether they were simply mistaken throughout the entire block of trials. As such, 31 of 1521 trials (2%) were removed for subsequent analyses.

Overall frequencies of detection reveal that CD gain manipulation is not very perceptible. On average, when a CD gain manipulation was presented, participants were able to detect it only 29.6% of the time. The frequency of detection increased with the level of stickiness as shown in Figure 7. At the highest level of stickiness, when motor space was expanded by a factor of 20, CD gain manipulation was only detected 82% of the time, and when motor space was expanded by a factor of 10, CD gain presence was only detected 65% of the time.

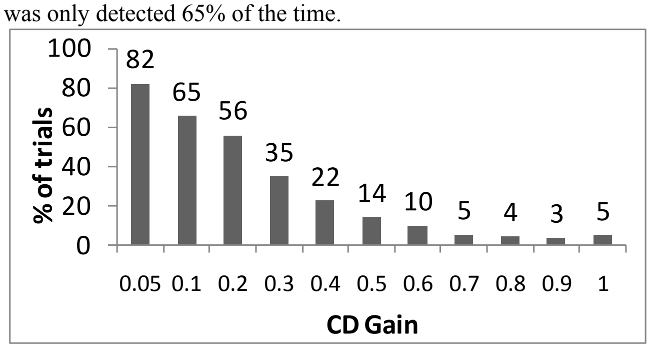

Figure 7. Percent of trials where CD gain was detected.

To determine whether participants were detecting CD gain changes successfully, we compared their performance to a 'random guess'. We compared participant detection rates to the random guess results in one-sample t-tests for each level of gain.

For CD gain values of 0.9 through 0.3, participants performed significantly worse than the random strategy (all p<0.001) – indicating that people had great difficulty noticing the change. For CD gain of 0.2, detection levels (57%) were not different than the random strategy (p>0.05). For gain of 0.1 and 0.05 – equivalent to target width expansions of 10x and 20x – participants detected the effect better than the random strategy (both p<0.001). In the distracter trials where there was no change in CD gain, users erroneously 'detected' change in 5% of the trials.

We were also interested in whether CD gain detection rates differed for varying target widths and amplitudes. A Univariate ANOVA with CD, A, and W as factors revealed no systematic differences in detection rates as a function of either amplitude ($F_{2,24}$=0.67, p>0.05, η^2=.05) or width ($F_{2,24}$=0.38, p>0.05,η^2=.03).

4.4 CD Gain Utility: Velocity-based Measures

Although MT results show that sticky targets help users to aim faster, velocity-based measures can help determine *why* sticky targets decrease MT (see section 5.1.2). Results from the repeated measures MANOVA revealed significant effects of CD gain manipulation on both peak velocity (PkV- $F_{5.8,69.9}$=63.0, p<0.001, η^2=0.84) and percent time after peak velocity (%afterPvK-$F_{6.8,81.2}$=57.5, p<0.001, η^2=0.83). Figure 8 shows how %after PkV increases with increasing CD gain, as predicted by the expansion of target width in motor space [7], and how PkV increases with decreasing gains, a result not expected if the benefits of target CD gain are only useful when the cursor is positioned above the target or if the benefits of targets CD gain manipulation are only explained by target width expansion in motor space. See Figure 9 for pairwise comparisons.

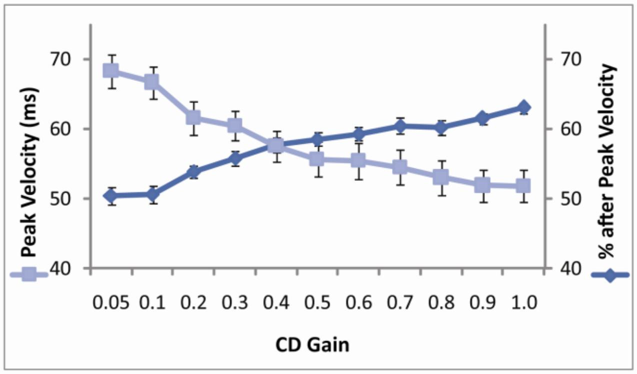

Figure 8. Mean peak velocity and % time after peak velocity ± SE.

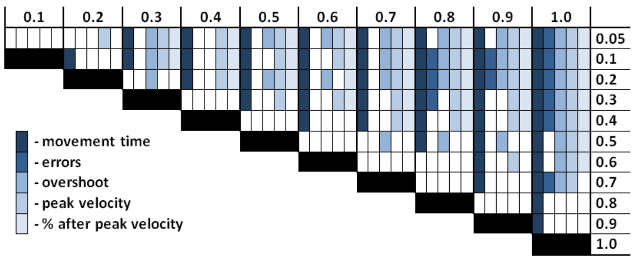

Figure 9. Pairwise comparisons of CD gain levels for each DV. All significant differences (shaded) are p<0.05 (Bonferroni adjusted). Measures (left to right) are movement time, errors, overshoots, peak velocity, and % time after peak velocity.

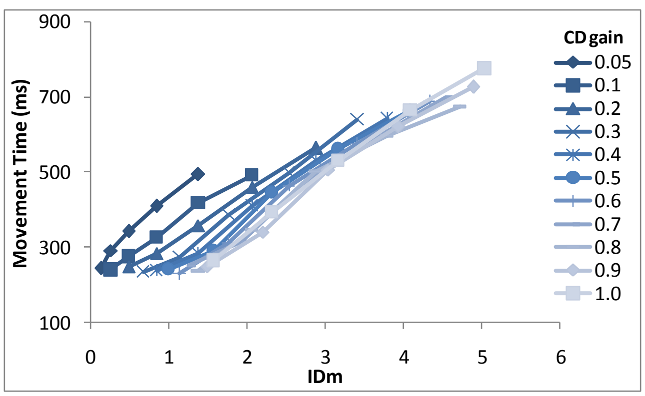

Figure 10. Mean MT as a function of IDm and CD gain.

4.5 CD Gain Utility: Modeling MT with Fitts's Law

Blanch et al. [7] investigated a doubling and quadrupling of target width in movement space (Wm), equivalent to CD gains of 0.5 and 0.25 in our study. They found that modeling movement time by calculating ID in motor space (IDm) provided a better linear fit than calculating ID in visual space (IDv). They also found that the full MT benefits of quadrupling motor space were not achieved due to the accompanying increased accuracy in that condition.

We performed linear regressions on mean MT using both IDv and IDm predictors (note that ID=log2(A/W+1) was used in the regression analysis to avoid negative IDs). Although we also achieved a slightly better linear fit using IDm, the r2 values of each approach are similar (IDv r2=0.84, IDm r2=0.87). Figure 10 shows how actual MTs for very sticky targets (low CG gains) are higher than predicted. As in [7], the full benefits of target width expansion in motor space are not being used. Although Figure 5 shows how accuracy improves with lower CD gain, error rate was not different for CD gains less than 0.8; thus the discrepancy is not fully explained by a speed-accuracy trade-off.

5 Discussion

There are four main results of our study:

R1: As expected from previous work, sticky targets are useful, and can provide substantial performance improvements;
R2: CD gain manipulation reduces targeting time significantly even with the smallest gain difference used in our study (0.9);
R3: Sticky targets are not easy to perceive, even when participants can compare before and after CD gain changes;
R4: There is a large gap between utility and perceptibility: sticky targets assist performance at low levels of stickiness, but only high levels are consistently noticed.

In the following sections, we discuss possible explanations for our results, ways that the findings can be generalized to real-world situations, and implications for design and further research.

5.1 Explanations for Results

5.1.1 Why Did People Not Notice Stickiness?

Our study showed that even high levels of CD gain were not consistently noticed. There are several possible reasons for this result. First, the effect of small changes to CD gain is subtle, and people may simply not be good at noticing these types of changes in the relative movement of the cursor. Second, it is possible that there is a degree of change blindness occurring, in that participants were focusing on the targeting task rather than the CD gain, and thus were looking primarily at the location of the cursor, not its movement characteristics (that is, the movement task obscured the perception task). Third, in a related phenomenon, called mindsight, observers of a changing stimulus may reliably sense a change, but not have a visual experience of it [26]. The subjective difference between *sensing* and *seeing* visual stimuli could extend to pseudo-haptic targets where users may *sense* CD gain changes without *feeling* them.

5.1.2 Why Does Stickiness Help?

There may be multiple valid explanations for why sticky cursors assist with aiming time. First, faster movement times from lower CD gain on targets could be due to an expansion of the target width in motor space [7]. The improved prediction of MT when using IDm rather than IDv supports this explanation. However, based on [21], if sticky targets *only* help because of increased target width in motor space, there should be an effect of CD gain on the percent of time spent decelerating, but not on peak velocity. Our results show improvements in both PkV and %after PkV with increased stickiness. Since our experiment did not decrease target amplitude (it increased effective A for the portion of the movement over the sticky target), there must be another explanation for the increases in peak velocity.

In Keyson's [18] studies, the ID of the task was held constant regardless of the use of CD gain adaptation, yet improvements in MT were seen with sticky targets. Supported by other experimental conditions where tactile feedback (through a force-feedback trackball) improved aiming times over CD gain manipulations alone, Keyson suggests an explanation called 'cursor catching' – where the user essentially 'brakes' over the target location. One of the criticisms of sticky targets is that unlike visual target expansion or vortex approaches where target amplitude is decreased, sticky targets only help users aim when they are over the target and not when they are en route to the target [3]. As a result of our pilot studies, we felt there might be improvements both overtop of and in between sticky targets, and wanted to explore evidence for Keyson's braking hypothesis.

In physical aiming to targets, impact with the target leads to a force opposite to the direction of movement, helping to decelerate the limb [31]. In [1], users slid a pen over a tablet to a target and either stopped (discrete task) or returned to the start position (reciprocal task), with or without a mechanical stop at the target. The authors found that movement times were shorter with the mechanical stop, and that users were moving faster (higher PkV) and spending less percentage of the movement decelerating (%after PkV). Although we do not actually provide a mechanical stop, users could also be aided by the sensation of ‘braking’ provided by the pseudo-haptic stickiness, which would be reflected in both peak velocity and %time after peak velocity. Our results showed changes in both of these measures, lending credence to Keyson’s braking explanation, and possibly explaining why our participants were moving at higher velocities without any changes to effective target amplitude.

Detracting from the braking explanation is that our participants only noticed CD gain manipulations reliably at gains less than 0.2, but we saw improvements in aiming time, starting with a CD gain of 0.9. If users were not detecting CD gain manipulations, how could they experience the pseudo-haptic sensation of stickiness necessary for the braking hypothesis? There is evidence that humans treat visual information differently depending on whether the task is related to perception or action, and that separate anatomical visual pathways support the two activities (see [6]). As such, we can use visual information to help us aim even when we do not consciously perceive it [6]. It is possible that participants used visual cues of cursor slowdown (stickiness) to help them decelerate without being consciously aware of it. In addition, similar to mindsight [26], it is possible that users simply *sensed* the changes in CD gain without *feeling* the changes.

5.1.3 Why is the Full Benefit of Expansion Not Seen?

Although the expanded targets created by CD gain resulted in improved performance, the improvement fell short of the equivalent conditions of ‘actual’ target width (where both motor and visual space are expanded). This phenomenon led to the poor correlation between our empirical data and Fitts’s law (section 4.5). We believe that this may have occurred because the targets’ visual size caused participants to underestimate the motor size. This is understandable, since participants did not know the target widths (in motor space), and could only discover these widths through experience. Although we used only the second block of trials in our analyses, the phenomenon may have persisted.

5.2 Generalizing the Results

5.2.1 Perceptibility Results

Our study tasks were designed so that CD changes would be maximally noticeable: the tasks involved repeated back-and-forth selections that provided participants with a baseline against which to judge differences, and participants did not do higher-level tasks that could distract them from noticing CD changes.

Real-world targeting actions are far more likely to be individual events that are mixed in with other actions, rather than repeated trials as seen in our study. When users experience the effect of CD gain, therefore, they will not have a clear memory of what targeting feels like without the effect. In addition, they will be engaged in a domain task that may further distract them from the change in CD gain. Therefore we believe that for single-target tasks in the real world, the perceptibility of CD gain should be no greater, and likely even less, than what we saw in the study.

5.2.2 Utility Results

Our study’s results about the utility of sticky targets can be generalized to other tasks and situations, because the basic targeting action used in all settings (both our experiment and real-world tasks) is highly similar. However, targeting actions in real-world tasks can be considered to be somewhat more difficult than those used in our study: they will in general be slower than in the study (since they are not repeated), and the targets will most likely be two- instead of one-dimensional. Nevertheless, the performance benefits of motor-space expansion are very likely to transfer to real-world tasks – at least in single-target situations.

5.2.3 Generalization to Multiple Targets

Our study provides evidence about performance and perceptibility in single-target situations only. In order for sticky targets to be used reliably by designers of everyday interfaces, we also need to understand how the effect will be experienced in multiple-target situations as well. There are three issues to be considered:

Do performance results change in multiple-target situations? Performance results could change for two reasons. First, en route targets increase the effective amplitude of the final target (since each of the en-route targets has an increased effective width). Given the difficulty of modeling empirical data with either IDv or IDm in the present study, it is currently unclear whether targeting in the presence of en route targets can be modeled using either visual amplitude or motor amplitude. Second, targeting performance could be reduced because of the disruption of cursor motion in the initial phase of targeting. It is not currently known how changes to CD gain affect the closed-loop phase of aiming.

Does perceptibility change due to the presence of 'en route' targets? The effects of CD gain in our study were not highly noticeable; however, the presence of sticky regions en route to the real target may make these effects more obvious. Interactions are likely with the size of the en route regions, the amount of gain in these regions, and the location of the en route regions in relation to the targeting movement (open- or closed-loop phase).

Do 'en route' targets cause frustration or errors? The results of the current study suggest that sticky targets could provide benefits without a large perceptual cost to the user; however, further work is needed to determine other user costs such as frustration, annoyance, and errors. One main issue to investigate is how targets are usually organized in current interfaces: for example, tight clustering of targets in UI elements such as toolbars is more common in applications than targets spread across a large region (as seen with icons on a desktop). Clustered targets are less likely to pose problems for sticky targets, because there are only a few regions with different CD ratios.

5.3 Implications for Design and Further Research

Our results have implications for the design of visual interfaces. These issues require further study, but we present them here to state the potential for our results in real-world systems.

Small amounts of stickiness could substantially improve performance without great cost. The gap between utility and perceptibility of sticky targets could present a major opportunity for designers. If further testing in realistic situations confirms our results, sticky targets could allow designers to improve performance in pointing-based visual interfaces, without causing any major changes to the user's experience. CD gains starting at 0.9 (corresponding to motor-space expansions of only 1.11 times the visual size of the target) provided significant benefit but were not reliably detectable by participants until CD gain of 0.1 (a motor-space expansion of 10x). For example, a gain of 0.7 (1.43 times width) led to a 12% reduction in targeting time, but was noticed by only 5% of participants.

Low CD gain may make targets too sticky. Although not significant, movement times and error rates show an upward trend at high levels of stickiness. When CD gain was manipulated by a factor of five or more, participants may have experienced some difficulty in entering the target in order to select it. If very sticky targets are used, system designers might consider leaving a few-pixel border around the inside of the target that is not as sticky.

CD gain may not be an effective way to simulate haptic feedback on targets. Because of the poor perceptibility of CD gain in our study, the effect may not be a good choice for adding pseudo-haptic feedback to targets. In these systems, the goal is not improvement of performance, but the provision of an experience that simulates tactile feedback. For this experience to be a successful simulation, however, the feedback must be perceivable, and our study suggests that low to moderate changes of gain are not easily perceived, at least in targeting situations. CD gain has been used as effective pseudo-haptic feedback in free-exploration tasks [20]; it may be that these settings allow the user to perceive the changes in gain more successfully. CD gain may still be an important element of a pseudo-haptic targeting interface, but it may need to be coupled with other more noticeable types of feedback such as auditory information.

6 CONCLUSION

Control-to-display ratio has been suggested as a mechanism for improving targeting performance, but little is known about how users experience these manipulations. Depending on how users perceive the effect, and how the perceptibility relates to utility, the experience could range from that of an unhelpful annoyance to a benefit that comes more or less for free. We carried out a study to explore the gap between the perceptibility and utility of CD gain in sticky targets. We found that there is a gap: even small changes to CD gain significantly improve targeting time, but users do not notice effects until the gain difference is large. Our results provide new understanding of how CD gain changes are experienced, and suggest that sticky targets have considerable potential for providing performance benefits without high costs to users.

Our initial study also suggests several directions for future work. We plan to carry out a second study to test sticky targets in more realistic targeting situations involving visual environments and tasks that are more similar to those seen in the real world. Second, we will explore the issue of multiple targets: we will carry out further experiments to gather baseline data with multiple targets, and also investigate the type and frequency of multiple-target situations in real interfaces. Last, we plan to develop a dynamic CD gain system that uses cursor velocity to adjust gain depending on the velocity profile of the user's current movement.

7 REFERENCES

[1] Adam, J. J., van der Bruggen, D.P.W., & Bekkering, H. (1993). The control of discrete and reciprocal target-aiming responses: Evidence for the exploitation of mechanics. Human Move. Sc., 12, 353–364.

[2] Ahlstrom, D. (2005). Modeling and improving selection in cascading pull-down menus using Fitts' law, the steering law and force fields. CHI'05, 61-70.

[3] Ahlström, D., Hitz, M., & Leitner, G. (2006). An evaluation of sticky and force enhanced targets in multi-target situations, Nordic HCI'06, 58-67.

[4] Balakrishnan, R. (2004). "Beating" Fitts' Law: virtual enhancements for pointing facilitation. IJHCS, 61(6), 857-874.

[5] Baudisch, P., Cutrell, E., Hinckley, K., & Eversole, A. (2005). Snap-and-go: Helping users align objects without the modality of traditional snapping. CHI'05, 301-310.

[6] Binsted, G., Brownell, K., Vorontsova, Z., Heath, M., & Saucier, D. (2007). Visuomotor system uses target features unavailable to conscious awareness. PNAS, 104 (31), 12669-72.

[7] Blanch, R., Guiard, Y., & Beaudouin-Lafon, M. (2004). Semantic Pointing: Improving target acquisition with control-display ratio adaptation. CHI'04, 519-526.

[8] Casiez, G., Vogel, D., Balakrishnan, R. & Cockburn, A. (in press). The impact of control-display gain on user performance in pointing tasks. JHCI.

[9] Cockburn, A. & Firth, A. (2003). Improving the acquisition of small targets. People and Computers XVII. 181-196.

[10] Cockburn, A. & Brewster, S. (2005). Multimodal feedback for the acquisition of small targets. Ergonomics 48(9), 1129-1150.

[11] Cockburn, A. & Brock, P. (2006). Human on-line response to visual and motor target expansion. GI'06, 81-87.

[12] Fitts, P.M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. J.Exp. Psych. ,47, 381-391.

[13] Fitts, P.M. & Peterson, J.R. (1964). Information capacity of discrete motor responses. J.Exp. Psych., 67, 103-112.

[14] Gallace, A., Tan, H.Z., & Spence, C. (2006). The failure to detect tactile change: A tactile analogue of visual change blindness. *Psychonomic Bulletin & Review*, 13, 300-303.

[15] Guiard, Y., Blanch, R. & Beaudouin-Lafon, M. (2004). Object pointing: a complement to bitmap pointing in GUIs. GI'04, 9-16.

[16] Gutwin, C. (2002). Improving focus targeting in interactive fisheye views. CHI'02, 267-274.

[17] Hurst, A., Mankoff, J., Dey, A.K., & Hudson, S.E. (2007). Dirty Desktops: using a patina of magnetic mouse dust to make common interactor targets easier to select. UIST'07, 183-186.

[18] Keyson, D.V. (1997). Dynamic Control Gain and Tactile Feedback in the Capture of Cursor Movements. Ergonomics, 12, 1287-1298.

[19] Lécuyer, A, Coquillart, S, & Kheddar, A. (2000). Pseudo-Haptic Feedback: Can isometric input devices simulate force feedback? IEEE VR2000, 18–22.

[20] Lécuyer, A, Burkhardt, JM, & Etienne, L. (2004). Feeling bumps and holes without a haptic interface: the perception of pseudo-haptic textures. CHI'04, 239-246.

[21] MacKenzie, C.L., Marteniuk, R.G., Dugas, C., & Eickmeier, B. (1987). Three-dimensional movement trajectories in Fitts' task: Implications for motor control. QJEP, 39A, 629-647.

[22] MacKenzie, I.S. (1992). Fitts' law as a research and design tool in human-computer interaction. JHCI, 7, 91-139.

[23] McGuffin, M. & Balakrishnan., R. (2005). Fitts' Law and Expanding Targets: Experimental studies and designs for user interfaces, TOCHI, 12(4), 388-422.

[24] Meyer, D., Abrams, R., Kornblum, S., Wright, C., & Smith, J. (1988). Optimality in human motor performance: Ideal control of rapid aiming movements. Psych. Review, 95, 340-370.

[25] O'Regan, J.K., Rensink, R.A. & Clark, J.J. (1999). Change-blindness as a result of "mudsplashes". Nature, 398, 34.

[26] Rensink R.A. (2004). Visual sensing without seeing. Psychological Science, 15, 27-32.

[27] Rensink R.A, O'Regan J.K, & Clark J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. Psychological Science, 8, 368-373.

[28] Rodgers, M. E., Mandryk, R. & Inkpen, K. (2006). Smart sticky widgets: Pseudo-haptic enhancements for multi-monitor displays. Smart Graphics'06, 194-205.

[29] Sarkar, M. & Brown, M. (1992). Graphical fisheye views of graphs. CHI'92, 83-91.

[30] Simons, D., Franconeri, S., & Reimer, R. (2000). Change blindness in the absence of visual disruption. Perception, 29, 1143-1154.

[31] Teasdale, N., & Schmidt, R.A. (1991). Deceleration requirements and the control of pointing movements. J. Motor Beh., 23, 131–138.

[32] Vitevitch, M.S. (2003). Change Deafness: The inability to detect changes between two voices. *J.Exp. Psych.*, 29, 333-342.

[33] Woodworth, R.S. (1899). The accuracy of voluntary movements, Psych. Review Monograph, Supp., 31-114.

[34] Worden, A, Walker, N, Bharat, K, & Hudson, S. (1997). Making Computers Easier for Older Adults to Use: Area Cursors and Sticky Icons. CHI'97, 266-271.

[35] Zhai, S., Conversy, S., Beaudouin-Lafon, M. & Guiard, Y. (2003). Human on-line response to target expansion. CHI'03, 177-184.