Multiple Haptic Targets for Motion-Impaired Users

Faustina Hwang, Simeon Keates, Patrick Langdon, P John Clarkson
University of Cambridge, Department of Engineering
Trumpington Street, Cambridge, UK CB2 1PZ
+44 1223 766958
{fh220, lsk12, pml24, pjc10}@eng.cam.ac.uk

ABSTRACT

Although a number of studies have reported that force feedback gravity wells can improve performance in "point-and-click" tasks, there have been few studies addressing issues surrounding the use of gravity wells for multiple on-screen targets. This paper investigates the performance of users, both with and without motion-impairments, in a "point-and-click" task when an undesired haptic distractor is present. The importance of distractor location is studied explicitly. Results showed that gravity wells can still improve times and error rates, even on occasions when the cursor is pulled into a distractor. The greatest improvement is seen for the most impaired users. In addition to traditional measures such as time and errors, performance is studied in terms of measures of cursor movement along a path. Two cursor measures, angular distribution and temporal components, are proposed and their ability to explain performance differences is explored.

Keywords

pointing devices, force feedback, cursor trajectory

INTRODUCTION

For most computer users, pointing devices are essential tools for effective interaction with graphical user interfaces (GUIs). Mouse use can make up between 31% and 65% of computer operation time [8]. Effective use of pointing devices, however, requires precise motor control and dexterity. In situations where accurate cursor control is not possible, GUIs become inaccessible in many ways.

Motion-impaired computer users often have difficulty with accurate cursor control [15]. Symptoms such as tremor, spasm, muscle weakness, partial paralysis, or poor coordination can make standard pointing devices difficult, if not impossible, to use. Although alternative methods of interaction with GUIs exist (e.g. MouseKeys, voice activation, switch-based scanning), these methods may be inappropriate or simply too slow for effective interaction. Yet the need to make computers accessible to those with disabilities is increasingly important, as reflected in legislation such as Section 508 of the *1998 Workforce Investment Act* [17].

CHI 2003, April 5–10, 2003, Ft. Lauderdale, Florida, USA.

A number of studies have shown that haptic feedback, i.e. feedback that allows users to *feel* a computer interface, can improve user performance in various tasks involving cursor control [e.g. 1, 3]. In particular, force feedback gravity wells, i.e. attractive basins that pull the mouse and cursor to the center of a target, have been shown to improve performance in "point-and-click" tasks. Hasser *et al* [6] found that this type of force feedback, provided by a FEELit mouse, could improve targeting time and decrease errors. Oakley *et al* [13] reported a reduction in errors with the use of gravity wells implemented on a PHANToM. Keates *et al* [9] found that for motion-impaired users, gravity wells could improve the time required to complete a "point-and-click" task by as much as 50%.

In these studies, however, force feedback was enabled on a single target only. For the successful implementation of force feedback in a realistic interface, issues surrounding haptic effects for multiple on-screen targets must be addressed. With more than one gravity well enabled, a user's cursor may be captured by the gravity wells of undesired *distractors* as it travels toward a desired target. This has the potential to cancel out the benefits of the force feedback, possibly yielding poorer performance than in its complete absence.

There have been few studies investigating performance in the presence of multiple haptic targets. Dennerlein and Yang [4] found that even with multiple haptic distractors along the cursor trajectory, performance in "point-and-click" tasks was greatly improved over a condition with only visual feedback. Study participants most often just plowed through the distractors, but at a cost of increased user frustration and effort. In contrast, Oakley *et al* [14] reported an increase in time when static attractive forces were enabled on multiple targets. This condition was, at best, not optimal, and at worst, detrimental to performance and subjective satisfaction when compared to the purely visual condition. Langdon *et al* [11] reported a performance improvement for motion-impaired users that was similar across four sizes of gravity wells on adjacent targets.

Given the conflicting reports of these studies, arising possibly from different experimental setups, the use of multiple haptic targets is an area that requires further investigation. This paper studies the performance of users with a wide range of capabilities in “point-and-click” tasks when multiple on-screen targets are haptically enabled, building on the previous literature in primarily two ways.

First, previous studies have most often evaluated performance in terms of task completion time and errors. However, differences in cursor movement between people, devices, and haptic conditions can exist in a variety of ways. Although traditional measures may show *that* a difference exists between conditions, establishing *why* they exist is more likely to be accomplished by analyzing the path of movement *throughout* a trial [12]. An understanding of how cursor movement is affected by the presence of multiple haptic targets can aid in the design of interfaces that are better suited to user needs. This paper proposes a new measure of cursor movement and illustrates how it may be applied to explain performance differences among haptic conditions, and consequently to suggest improvements for interface design. Second, there have been no explicit studies of the importance of distractor location relative to the target when multiple targets are haptically enabled. An understanding of how the spatial relationship between haptic targets affects performance will facilitate the design of improved layouts.

MEASURES OF CURSOR MOVEMENT

MacKenzie *et al* [12] proposed seven new accuracy measures for evaluating computer pointing devices. Most of these measures were based on the assumption that a “perfect” trajectory is one that follows the task axis, defined as the straight line connecting the start and end points of the task. However, as the instantaneous task axis can vary throughout the cursor trajectory, Keates *et al* [10] proposed six measures that were not dependent on the initial task axis. In this section, two measures that complement those previously proposed are described.

Angular Distribution (*AD*)

A new cursor measure is proposed that captures the angular distribution (*AD*) of a cursor trajectory. The task axis is defined to be the line from the start point to the end point of a defined task. Taking the task axis to be at 0 degrees, *AD* counts the frequency of cursor samples occurring in 15 degree segments about the target. In other words, if a cursor path is composed of sample points p_i, each forming an angle θ_i with the task axis (Figure 1), then *AD* is a histogram of all the θ_i values in a trajectory.

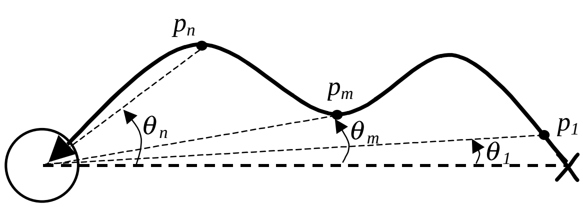

Figure 1. Sample points along a cursor path.

Figure 2 shows sample cursor traces from two motion-impaired computer users performing a “point-and-click” task. Although both users have Cerebral Palsy, the extent to which their motion is impaired is quite different, with user HA1 exhibiting only mild impairment in the dominant hand and arm, and IP1 exhibiting severe motor control difficulties in the dominant hand and arm. Plots of *AD* for these two traces are shown in Figure 3.

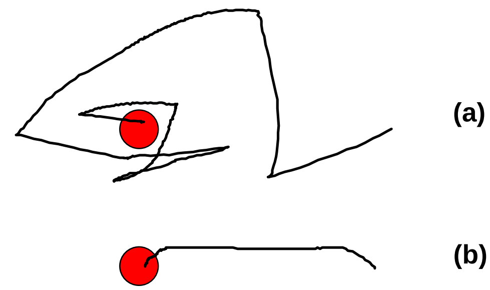

Figure 2. Sample cursor traces for (a) IP1 and (b) HA1.

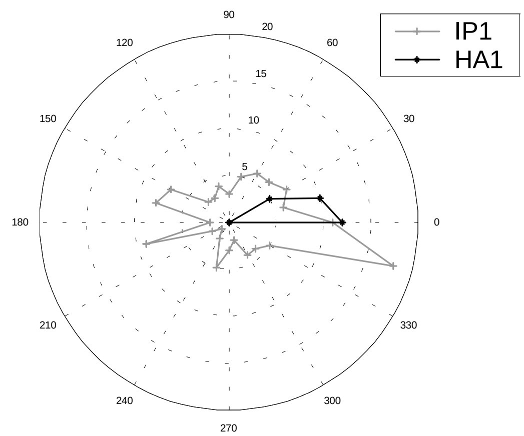

Figure 3. Plots of *AD* for the cursor traces in Figure 2.

The area contained within the *AD* plot conveys information about how far the cursor travels during a single trial. These plots reflect the fact that the cursor of IP1 travels a greater distance than that of HA1 to complete the same task. Also, *AD* for HA1 lies entirely between 0 and 30 degrees, reflecting a movement that deviates little from the task axis. For IP1, however, the movement is much less direct. *AD* shows that this user is just as likely to be approaching the target from 0 degrees as from 195 degrees. This gives further support for the development of cursor measures that do not assume a constant task axis, as it illustrates how greatly the instantaneous task axis can vary during a trial.

Temporal Components

A study of different components of the total task completion time can reveal insights to cursor behavior that are masked when a single “gross” measurement is used [2]. When two haptic targets are on the screen, the total time can be decomposed into three components based on the location of the cursor: (1) time spent inside the desired target, (2) time spent inside the distractor, and (3) time spent outside both haptic targets.

The potential of these two measures of cursor movement to provide explanations for performance differences were illustrated in a study of multiple haptic targets.

EXPERIMENT

An experiment was conducted to investigate the performance of computer users in "point-and-click" tasks when multiple on-screen targets are haptically enabled. Users performed the tasks under two force feedback conditions: (1) without force feedback, and (2) with gravity wells on both a desired and a distractor target. The distractor position relative to the task axis was also varied (see Figure 6).

The hypotheses were as follows:

H1: Force-feedback gravity wells would significantly improve task completion times and errors when compared with the no-force feedback condition.

H2: The distractor location relative to the cursor's start position and the target location would have a significant effect on task completion times when the force feedback is on, but not when the force feedback is inactive.

H3: There would be significant differences in times and errors between users with different levels of physical capability.

Participants

Ten volunteers with motion-impairments and eight without impairments participated in the study. The group with motion-impairments represented a wide range of capabilities, exhibiting symptoms including spasm, tremor, coordination difficulties, stiffness, and weakness in the dominant hand and arm. The users were affected by Cerebral Palsy (7), Friedrich's Ataxia (1), head injury (1), and spinal cord injury (1). Although the users had predominantly the same clinical diagnosis, Cerebral Palsy, the severity of the impairment ranged from mild to severe. Three users performed "point-and-click" tasks with two hands, using one hand to navigate the mouse and the other to press a button for selection.

Task

The task was a multi-directional point-and-click task, using a Logitech Wingman force feedback mouse for input. This device, shown in Figure 4, is capable of generating a wide range of haptic effects, including vibrotactile sensations and directional forces.

Figure 4. The Logitech Wingman force feedback mouse.

Potential target positions, indicated by faint gray circles on the screen, were located at the center and vertices of a regular hexagon with an edge length of 250 pixels (Figure 5). The center circle was initially filled in red, and users moved the cursor inside the red target and clicked the left mouse button to select it. Once selected, the red target appeared in a new position, along with a distractor. The distractor, shown as a white circle drawn with a black line, appeared in one of four locations relative to the task axis. Taking the task axis to be 0 degrees, distractors were located at 180, 90, 0, and -90 degrees, always 80 pixels from the target's center (Figure 6). After every selection, the target appeared in a new position, randomly selected from the set of adjacent positions. Targets and distractors were 40 pixels in diameter, and the screen resolution was 1024 by 768 pixels.

A trial was defined to be one complete target selection. The time to complete each trial, the number of mouse clicks in a trial, and the cursor position throughout the trial were all automatically logged after the first target had been selected. Data collection was then continuous for 16 subsequent trials, so the end of one selection became the beginning of the next. Breaks were taken between each block of 16 trials.

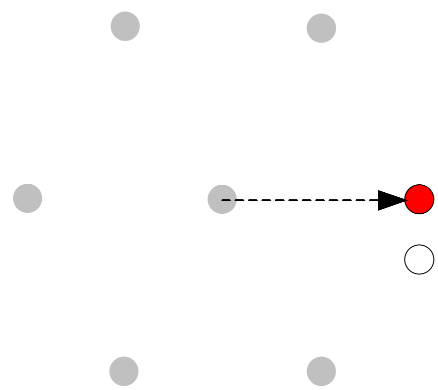

Figure 5. Target arrangement in the "point-and-click" task.

Design

The experiment was a 2x4 factorial within-subjects design. Each user completed the "point-and-click" tasks with the force feedback *off* and *on*. When the force feedback was *off*, the Wingman operated as an ordinary mouse. With it *on*, both the target and the distractor sat at the center of a circular gravity well with a radius two times that of the target. When the cursor entered the gravity well, a spring force pulled the mouse toward the center of the target. The distractor location was also varied, appearing at 180, 90, 0, or -90 degrees relative to the task axis (Figure 6).

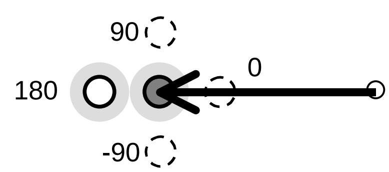

Figure 6. Four possible distractor locations and the extent of the gravity wells.

Within a block of sixteen continuous trials, the force feedback condition was held constant while each distractor location was presented four times. The order of the locations was determined using a method of random selection without replacement. The order of the force feedback levels between blocks was counterbalanced.

RESULTS

Of the study participants, user IP1 experienced the most difficulty performing the "point-and-click" tasks, reflected in task times, without force feedback, that were almost twelve times longer on average than those of the other users, and with a much greater variance. A separate analysis for this user was therefore more appropriate.

Overall

A two-factor repeated measures ANOVA with one between-subjects factor was performed on the mean task completion times and errors for all users except IP1. The error count for one trial is defined as the number of mouse clicks occurring outside the target during the trial. The between-subjects factor was *impairment*, with two levels: motion-impaired (*MI*) and able-bodied (*AB*).

Force-feedback gravity wells gave an overall improvement in times ($F_{1,15}$ = 13.316, p = 0.002). The improvement for motion-impaired users was 33% (= 1s) compared with the 10% (= 0.1s) observed for able-bodied users (Figure 7).

There was also an effect of distractor location on times, but only when the force feedback was *on* ($F_{3,45}$ = 11.244, p < 0.001). With the force feedback *on*, the effect was due to a difference between the distractor at 0 degrees and the other three locations. When the distractor lay directly between the start point and the target, times were significantly slower (Post-hoc comparisons with Bonferroni adjustments, all p<0.003). For able-bodied users, the increase in time was 0.15s (=15%), thereby not only canceling out any benefits of having force feedback on the target, but also hindering performance and yielding poorer times than in the complete absence of force feedback. In contrast, for the motion-impaired users, although the distractor slowed times by 0.4s (=17%), an improvement was still gained over the force feedback *off* condition.

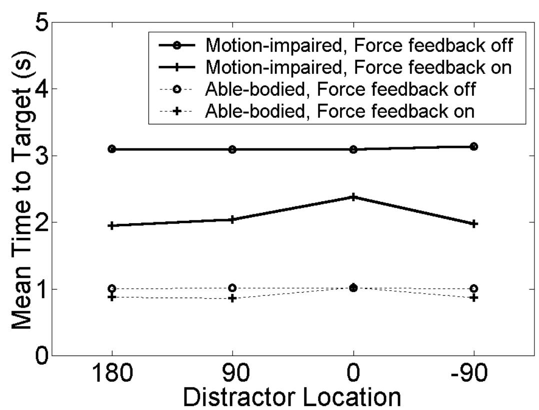

Figure 7. Interaction diagram for times (all users bar IP1).

When performance was measured in terms of errors, there was no significant difference between motion-impaired and able-bodied users (Figure 8). Overall, force-feedback gravity wells improved error rates by 40% ($F_{1,15}$ = 29.034, p < 0.001), reducing errors by 0.19 mouse clicks per trial. In contrast with the results for task completion times, distractor location did not have a significant main effect on errors. This means that although having a distractor located between the start point and the target was detrimental in terms of times for able-bodied users, a benefit was still gained in terms of error reduction.

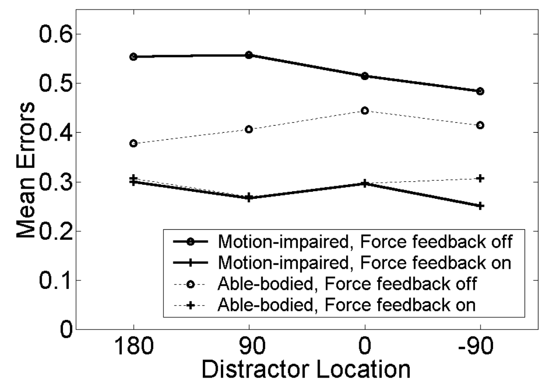

Figure 8. Interaction diagram for errors (all users bar IP1).

Participant IP1

A two-factor ANOVA was performed on times and errors for IP1. The mean times are shown with 95% confidence intervals in Figure 9. Force feedback gravity wells improved times by 45% ($F_{1,192}$ = 16.060, p < 0.001), equivalent to an 18 second reduction in time. In absolute terms, this is 18 times greater than the improvement observed for the other motion-impaired users. This result is in accordance with those reported in [9], providing further evidence that haptic feedback may be of greatest benefit to those who require the most assistance. In contrast to the overall result for the other users, for IP1, times were consistent for all four distractor locations ($F_{3,192}$ = 2.257, p = 0.083).

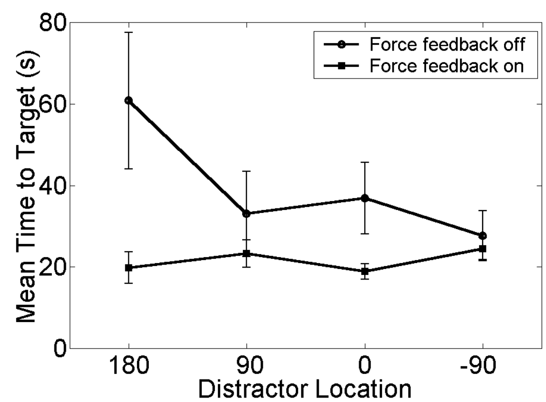

Figure 9. Mean times and 95% confidence intervals (IP1).

Results in terms of errors were similar (Figure 10). Force feedback improved errors by 51% ($F_{1,192}$ = 10.619, p =

0.001), a reduction of 0.62 mouse clicks per trial, and the result was unaffected by distractor location.

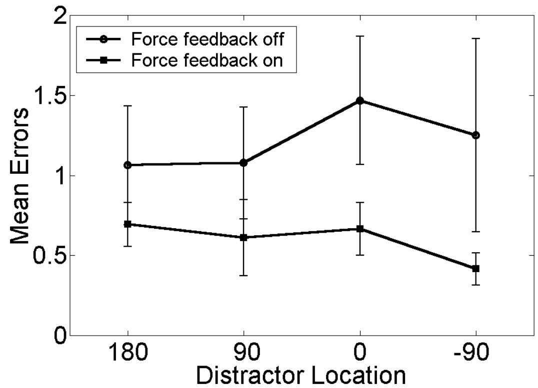

Figure 10. Mean errors and 95% confidence intervals (IP1).

CURSOR TRAJECTORY ANALYSIS

Although the measures of time and errors have shown that there are differences between conditions and users, they give little explanation about *why* these differences occur. Using participants HA1 and IP1 as illustrative examples, this section discusses the potential for *AD* and temporal components analysis to differentiate between two very different behaviors, as well as provide a possible explanation for (1) why times increased for the 0 degrees distractor location for HA1, in contrast with (2) why the distractor location had no effect for IP1, and (3) how the force feedback gravity wells are assisting users.

Figure 11 shows sample cursor traces for HA1 and IP1 when the force feedback is *on* and the distractor is at 0 degrees. These traces can be compared with those of Figure 2, recorded when the force feedback was *off*. The *AD*s averaged over all the cursor traces recorded for distractors at 90 and 0 degrees are shown in Figure 13. The plots indicate that the cursor of IP1 travels much further than that of HA1 on average, which may help explain IP1's much higher task completion times.

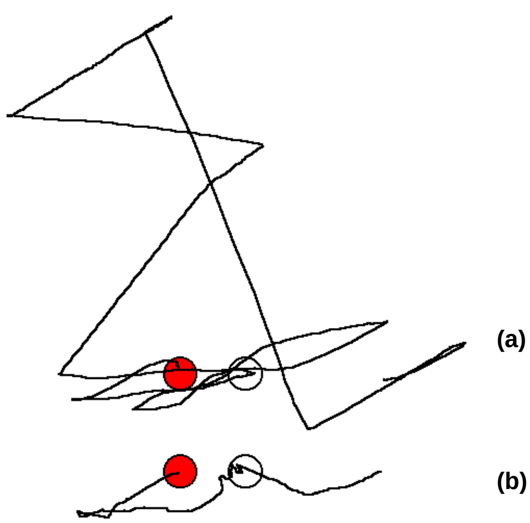

Figure 11. Cursor traces for (a) IP1 and (b) HA1.

To help explain why times increased for the 0 degree position, the *AD* for user HA1 across the distractor locations is compared (Figures 13a and 13c). The plots indicate that cursor traces for HA1 are predominantly directly toward the target, with movement constrained to within 15 degrees of the task axis. This is true even when the distractor is located directly in front of the target. That is, this user continued to navigate close to the task axis, rather than, for example, diverting widely *around* the distractor to avoid it completely. The cursor is then pulled into the distractor, and captured there momentarily as the gravity well exerts forces that resist movements to exit the well. The result is an increase in the amount of time spent inside the distractor, and consequently, an increase in the total task completion time. In fact, the extra time spent in the distractor accounts almost entirely for the increase in total time. This is reflected in Figure 12a which shows the total movement time (*Total*) broken down into three components: (1) time spent inside the desired target (*InTarg)*, (2) time spent inside the distractor (*InDist*), and (3) time spent outside both haptic targets (*Other*).

The plot of *AD* also offers a further explanation for the time increase. When the distractor is at 0 degrees, *AD* shows increased cursor movement near 180 degrees (Figure 13c). This shows that the measure is sensitive to the fact that, in escaping the forces of the distractor's gravity well, the cursor of HA1 often overshoots the target. Actions to correct this overshoot incur a slight time penalty.

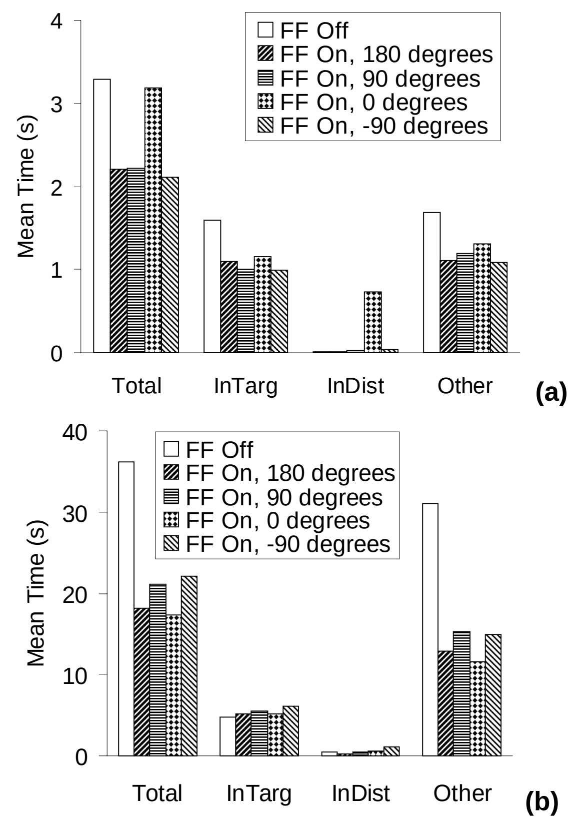

Figure 12. Time components for (a) HA1 and (b) IP1.

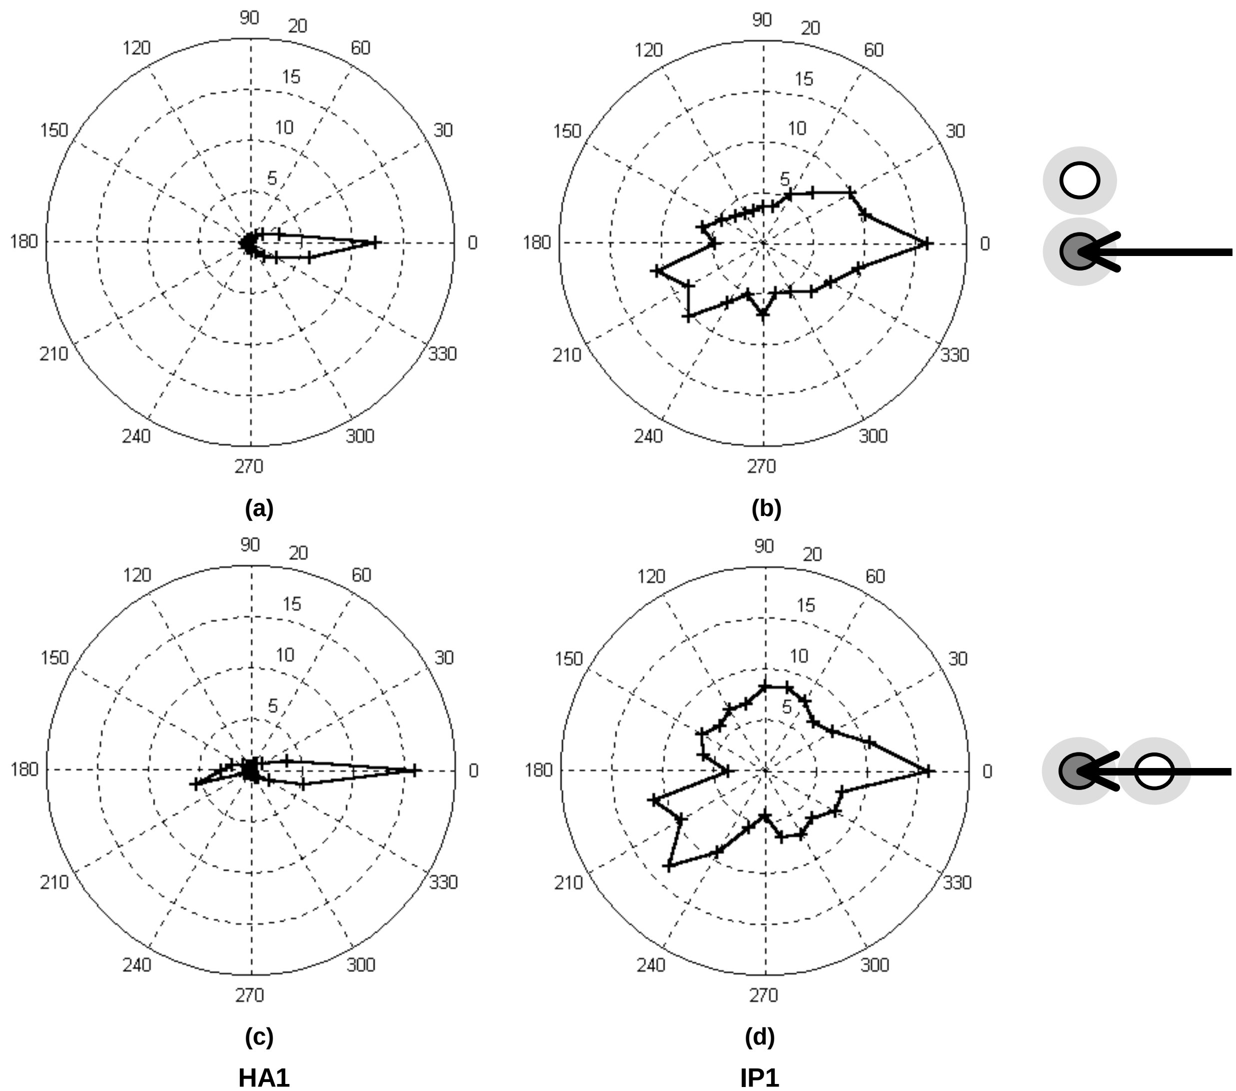

Figure 13. Plots of *AD* for (a) HA1 with the distractor located at 90 degrees, (b) IP1 with the distractor located at 90 degrees, (c) HA1 with the distractor located at 0 degrees, and (d) IP1 with the distractor located at 0 degrees.

To help explain why the distractor location had no effect for IP1, the plots of *AD* are again compared across distractor locations (Figures 13b and 13d). The cursor traces of IP1 exhibit a much wider distribution of angles than those observed for HA1. Within a single trial, IP1 can be equally likely to be approaching the target from the task axis at 0 degrees, or from 225 degrees (see Figure 13d). Consequently, the initial task axis and the relative location of the distractor bear little relevance. Having the distractor lie on the task axis does not increase the likelihood of IP1 getting pulled into it. This is reflected in the temporal components analysis, which shows similar amounts of time spent inside the distractor for all four distractor locations (Figure 12b).

Although force feedback proved to be of benefit for both HA1 and IP1, a study of the temporal components reveals that the gravity wells assist the two users in different ways. For HA1, the time savings comes from a reduction in the amount of time spent outside the two targets, as well as a reduction in the amount of time spent inside the desired target. The former is likely due to an increase in cursor speed once the cursor enters the gravity well surrounding the target. The latter is likely due to a faster reaction time between entering the target and clicking, possibly suggesting that, for this user, using haptic feedback to signal the arrival of the cursor at the target's centre is faster and more effective than using just visual feedback.

In contrast, the majority of time saved for IP1 has been due entirely to a reduction in the amount of time spent outside the two targets. The haptic feedback does not assist IP1 to click sooner after the cursor has entered the target, but the primary benefit is in helping the user to reach the target and to stay inside it.

DISCUSSION

The results of this study provide further evidence to support the addition of force feedback gravity wells to GUIs as a method of assisting users in performing “point-and-click” tasks. One major concern with the presence of multiple gravity wells was that users would be subjected to undesirable forces as the cursor passed over undesirable targets, consequently hindering progress and yielding poorer performance than in the complete absence of force feedback. However**,** this study showed that this simple haptic effect still improved performance over the non-haptic condition by reducing average times and errors, even in the presence of a gravity well on an adjacent distractor.

The present study has built on the literature by explicitly studying the effect of distractor location on performance. When the force feedback was *off*, no difference was found among the four distractor locations. When the force feedback was *on*, distractor location became significant for most study participants in the following ways.

Times were slower when the distractor was located between the cursor’s start point and the target, compared with other distractor positions. This can be attributed to the capture of the cursor by the distractor’s gravity well and extra time spent inside the distractor. In terms of haptic interface design, this arrangement is less than optimal, and a different target layout may be more appropriate, particularly for more able-bodied users. It is worth noting, however, that for motion-impaired users, times and errors were still improved over the force feedback *off* condition. Thus, the benefits of having the gravity well on the target outweighed the cost of having the cursor captured by the distractor. This is encouraging in terms of interface design, since in an actual interface with limited screen space, it may be difficult to avoid having potential targets aligned along the path of movement. Whether or not this result holds true when multiple distractors are present, however, is a topic currently under investigation.

Similar completion times were observed across the three conditions where the distractors were *not* located along the straight-line path to the target. In contrast with the 0 degree condition, users did not spend any more time inside the distractor than in the non-haptic condition. The implication for interface design is again encouraging, as the results suggest that similar time improvements can be gained in a more realistic situation where many targets are aligned in a row perpendicular to the cursor’s direction of travel.

This paper has also illustrated the potential of new measures to extract cursor movement characteristics that are not captured by the traditional measures of time and errors. By considering the path of movement, these new measures have the potential to explain *why* differences exist among users, to highlight difficulties that can occur, and therefore to provide insight into appropriate methods of assistance. For example, using a measure of angular distribution, *AD*, it was possible to identify that IP1 experienced particular difficulty in maintaining a consistent angle of approach to the target, and consequently the location of the distractor relative to the task axis had no effect. This user might benefit from a different form of haptic assistance such as tunnels [3, 7], which can help reduce the amount of angular deviation in the cursor path. The measure *AD* was able to show that the needs of IP1 were distinctly different from those of HA1 who most often moved to the target along the task axis, and consequently was frequently pulled into the distractor when the distractor was positioned in front of the target. This user would likely not gain much benefit from a haptic tunnel, but would benefit more from an alternative target layout.

Although these are only two examples, they are indicative of the range of differences that can exist among potential users and their needs, as well as of the range of methods of assistance that users may require. Computer interfaces need to be able to accommodate this diversity, and provide assistance that is appropriate to each individual’s ability.

In an environment where computers are often shared, or in a situation where a person’s ability changes over time [5], configuring an interface for each user can be difficult and time consuming [16]. People with motion-impairments can also be faced with an added difficulty in that the process of making an interface physically accessible first requires physical access of the interface. In these situations, a perceptive user interface that could automatically characterize a user’s behavior and implement an appropriate target layout and method of assistance would be more useful. In the same way that measures of typing have been used to identify keyboarding difficulties and to automatically adjust keyboard settings [16], measures of cursor trajectories may be used to automatically choose target sizes, target spacing, and forms of haptic assistance that will be most effective for a particular user.

To accomplish this, a single cursor measure is clearly insufficient. A measure such as *AD* can capture just one aspect of cursor movement, and will have limited ability to consistently distinguish between users of different capabilities [10]. Rather, a complementary set of measures, including traditional measures such as time and errors, as well as more recently developed ones such as those proposed in [12] and [10], need to be combined. In this way, it may be possible to form a fundamental set that can completely characterize a user’s behavior.

FUTURE WORK

Once a user’s area of difficulty has been identified, it is then equally important to implement an effective method of alleviating or eliminating the difficulty. Haptic feedback appears to be a promising approach for providing assistance to users with a wide range of capabilities, especially given the affordability of the Logitech Wingman force feedback mouse as an input device. Although this study has shown that force feedback gravity wells are beneficial for “point-and-click” tasks in the presence of two on-screen targets, the performance of motion-impaired users interacting with more realistic interfaces involving multiple haptic distractors remains to be fully investigated. Other forms of haptic feedback, as well as other GUI interaction tasks must also be studied.

CONCLUSION

Force feedback gravity wells can be effective at improving times and error rates for computer users in “point-and-click” tasks, even in the presence of a competing gravity well on an adjacent distractor. For motion-impaired users, the benefit of having a gravity well on the desired target outweighs the cost of being pulled into the distractor. For most users, placing the distractor between the cursor’s start point and the target is less than optimal. Cursor measures that capture aspects of a cursor trajectory have been used to demonstrate that a study of cursor movement along the path can help explain performance differences. Cursor measures can also give insight into the wide range of difficulties that users may experience. In this way, they are valuable in the development of accessible interfaces capable of responding to the needs of users with a wide range of capabilities.

ACKNOWLEDGMENTS

The authors thank the volunteers and staff of the Papworth Trust. This work is partially funded by the Canadian Cambridge Trust, NSERC, and the EPSRC.

REFERENCES

1. Akamatsu, M., MacKenzie, I. S., and Hasbrouq, T. A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device, *Ergonomics 38* (1995), 816-827.
2. Akamatsu, M. and MacKenzie, I. S. Movement characteristics using a mouse with tactile and force feedback, *International Journal of Human-Computer Studies 45*, (1996), 483-493.
3. Dennerlein, J. T., Martin, D. B., and Hasser, C. Force-feedback improves performance for steering and combined steering-targeting tasks. In *Proceedings of CHI 2000* (The Hague, 2000), ACM Press, 423-429.
4. Dennerlein, J. T. and Yang, M. C. Haptic force-feedback devices for the office computer: performance and musculoskeletal loading issues, *Human Factors 43*, 2 (2001), 278-286.
5. Gregor, P., Newell, A. F., and Zajicek, M. Designing for dynamic diversity – interfaces for older people. In *Proceedings of ASSETS 2002* (Edinburgh UK, 2002), ACM Press, 151-156.
6. Hasser, C. and Goldenberg, A. User performance in a GUI pointing task with a low-cost force-feedback computer mouse. In *Proceedings of the ASME Dynamic Systems and Control Division* (1998), American Society of Mechanical Engineers, 151-156.
7. Hwang, F., Keates, S., Langdon, P., and Clarkson, P.J. Using cursor measures to determine appropriate computer input assistance for motion-impaired users. To appear in *Proceedings of the Annual Conference of the IEEE EMBS 2002* (Houston TX, 2002).
8. Johnson, P.W., Hewes, J., Dropkin, J., and Rempel, D.M. Office ergonomics: motion analysis of computer mouse usage. In *Proceedings of the American Industrial Hygiene Association* (Fairfax VA, 1993), 12-13.
9. Keates, S., Langdon, P., Clarkson, J., and Robinson, P. Investigating the use of force feedback for motion-impaired users. In *Proceedings of the 6th ERCIM Workshop* (Florence Italy 2000), 207-212.
10. Keates, S., Hwang, F., Langdon, P., Clarkson, P. J., and Robinson, P. Cursor measures for motion-impaired computer users. In *Proceedings of ASSETS 2002* (Edinburgh UK, 2002), ACM Press, 135-142.
11. Langdon, P., Hwang, F., Keates, S., Clarkson, P. J., and Robinson, P. Investigating haptic assistive interfaces for motion-impaired users: Force-channels and competitive attractive-basins. In *Proceedings of Eurohaptics 2002* (Edinburgh UK, 2002), 122-127.
12. MacKenzie, I. S., Kauppinen, T., and Silfverberg, M. Accuracy measures for evaluating computer pointing devices. In *Proceedings of CHI 2001* (Seattle WA, 2001), ACM Press, 9-15.
13. Oakley, I., McGee, M.R., Brewster, S.A. and Gray, P.D. Putting the feel in look and feel. In *Proceedings of CHI 2000* (The Hague, 2000), ACM Press, 415-422.
14. Oakley, I., Adams, A., Brewster, S.A. and Gray, P.D. Guidelines for the design of haptic widgets. In *Proceedings of BCS HCI 2002* (London UK, 2002), Springer, 195-212.
15. Trewin, S. and Pain, H. Keyboard and mouse errors due to motor disabilities, *International Journal of Human-Computer Studies 50*, 2 (1999), 109-144.
16. Trewin, S. Accelerating assessment with self-optimizing devices. In *Proceedings of ICCHP 2002* (Linz Austria, 2002), Springer, 273-275.
17. Workforce Investment Act of 1998, Section 508. US Public Law 105-220 (1998).