The Effects of Feedback on Targeting Performance in Visually Stressed Conditions

Julie Fraser and Carl Gutwin

Department of Computer Science, University of Saskatchewan
57 Campus Drive, Saskatoon, SK, S7N 4J2
+1 306 966-8646
[jdf804, gutwin]@cs.usask.ca

Abstract

In most graphical user interfaces, a substantial proportion of the user's interaction involves targeting screen objects with the mouse cursor. Targeting tasks with small targets are visually demanding, and can cause users difficulty in some circumstances. These circumstances can arise either if the user has a visual disability or if factors such as fatigue or glare diminish acuity. One way of reducing the perceptual demands of targeting is to add redundant feedback to the interface that indicates when the user has successfully acquired a target. Under optimal viewing conditions, such feedback has not significantly improved targeting performance. However, we hypothesized that targeting feedback would be more beneficial in a visually stressed situation. We carried out an experiment in which normally-sighted participants in a reduced-acuity environment carried out targeting tasks with a mouse. We found that people were able to select targets significantly faster when they were given targeting feedback, and that they made significantly fewer errors. People also greatly preferred interfaces with feedback to those with none. The results suggest that redundant targeting feedback can improve the usability of graphical interfaces for low-vision users, and also for normally-sighted users in visually stressed environments.

Key words: Extraordinary HCI, accessibility, low-vision users, targeting, redundant targeting feedback.

1 Introduction

The vast majority of current graphical user interfaces involve manipulation of onscreen artifacts with a mouse-controlled pointer (Johnson et al. 1995). The core activity in these manipulations, one that users carry out over and over again, is targeting—the act of moving the pointer onto a manipulable region of the screen such as a button, a window border, a selection handle, or a scrollbar arrow. For most users, targeting does not present many problems; however, for users in a visually stressed environment, targeting can be an arduous task. For example, a person with reduced visual perception may be unable to locate the pointer to begin with, may lose track of the pointer en route to the target, or may have difficulty determining that the pointer is correctly placed on the target.

In this paper, we investigate the problem of making targeting tasks easier for visually stressed users—in particular, users who have reduced visual acuity. Two main groups of users in this category are people with visual disabilities and the elderly, but it also includes ordinary users in settings where bad lighting, fuzzy displays, or fatigue reduce visual acuity in normally-sighted users. As Alan Newell (1995) has pointed out, many situations exist where ordinary users are artificially disabled by environmental factors, and considering the design of interfaces for people with visual disabilities can have unexpected benefits for all users.

Our approach is to simplify targeting by providing the user with additional feedback during the targeting process. In particular, we consider the usefulness of feedback that indicates when the pointer enters or leaves a target. Although this kind of targeting feedback not been found to cause significant improvements in visually optimal settings (Akamatsu, MacKenzie, and Hasbrouq 1995), we believe that it will have a greater effect in a visually stressed situation. To test the usefulness of assistive feedback, we carried out an experiment where users in visually stressed conditions carried out targeting tasks with and without targeting feedback. The experiment is described below. First, however, we briefly discuss three areas that underlie the research: targeting, low-vision users, and redundant feedback.

1.1 Targeting

Targeting is the act of pointing to and selecting an object on the screen (Baecker et al. 1995). All direct-manipulation actions in graphical interfaces begin with a targeting task: for example, pressing a button, selecting text in an editor, or choosing a menu item all begin with the same user action of moving and positioning the mouse pointer. When the pointing device in the interface has an on-screen pointer (as opposed to a touch-screen or a light pen), we can divide targeting into three distinct stages: locating, moving, and acquiring.

Locating is the act of finding the mouse pointer on the computer screen when its position is unknown. Moving is the act of bringing the pointer to the general vicinity of the target, requiring the user to stay aware of the pointer's position as it travels across the screen. Acquiring is the final phase, and is the act of precisely setting the pointer over the target and determining that the pointer is correctly positioned. Acquisition requires greater fine motor control and attention to visual detail, and is the phase that we concentrate on in this work.

According to Fitts' law, targeting difficulty is determined by the size of the target and its distance from the starting location (e.g. Mackenzie 1995). Acquisition, however, is primarily affected by target size. In a standard Windows environment, there are several interface elements that are small enough to become potential targeting problems. The smallest common targets are window border at four pixels wide, and window splitters and tab stop markers at six pixels wide (see Figure 1). Other small targets include object selection handles in drawing programs (seven pixels) and window close buttons (12 pixels). Icons may also appear to have a small selectable region depending upon the visible picture, although the actual area of an icon is generally a larger region around the picture.

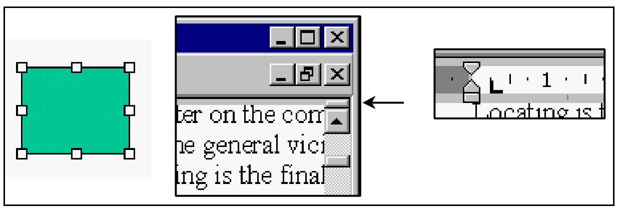

Figure 1. Small targets in the Windows environment. Left: object selection handles. Middle: window close buttons and window splitter (at arrow). Right: tab stop and indent markers.

1.2 Low-vision users

Low vision users are those people with a profoundly reduced degree of visual perception, but with enough usable eyesight to operate computer applications (Jacko and Sears 1998). There are a wide range of visual disabilities, but one of the main factors that affects people's use of graphical interfaces is visual acuity. A person with normal eyesight has a visual acuity of 20/20—from a distance of twenty feet they are able to see what any other person with normal vision can see from the same distance. However, a person with a visual acuity of 20/70 (for example) has significantly less eyesight. At twenty feet they are only able to see the level of detail that a person with normal vision sees from seventy feet. Acuity in the range of 20/70 to 20/160 constitutes moderate vision loss, while an acuity in the range of 20/160 to 20/400 constitutes severe vision loss (Levack 1994). We will use the term "low vision" to refer to both groups.

Although some assistive technologies exist for low vision users (e.g. screen magnification software), most users carry out their tasks with standard hardware and software, and most operate in a graphical interface environment (Fraser 1998). Graphical user interfaces are widely regarded as a significant step forward in the usability of computer systems. However, the shift from command line interfaces to graphical interfaces implies a trade-off for all types of computer users: while direct manipulation reduces the cognitive load placed on the user, it puts an increased demand on the perceptual systems. Graphical environments thus present a particular challenge for the sight impaired, and targeting is one of the major problems. The mouse cursor is a small, fast-moving object that can disappear against a non-contrasting background or can become lost in screen clutter. Target acquisition is particularly problematic because it requires precise visual discernment of pointer and target.

1.3 Redundant targeting feedback

Graphical user interfaces, by definition, provide a basic level of visual feedback to support targeting tasks—namely, the visual representations of the on-screen pointer and the target object. However, several applications have gone beyond this basic level to provide additional targeting feedback. This assistive information has been primarily visual, but can also be auditory or tactile information.

A common visual technique involves highlighting a selectable object when the mouse pointer enters the object's boundary. The highlight indicates that the pointer is correctly positioned to select the target. The technique can be seen in menus (see Figure 2) and more recently in application toolbars (see Figure 3). A second visual approach changes the appearance of the on-screen pointer rather than the appearance of the target. For example, when the pointer moves over a selectable window border in MS Windows, the cursor changes to indicate the resizing operation that can be performed (see Figure 4).

Auditory and tactile targeting feedback is less common in commercial applications, but has been used in assistive technology for the blind, and in a few research systems (Kline & Glinert 1995; Vanderheiden 1989). For example, Kline and Glinert explored the application of feedback intended specifically for the partially sighted. They used sounds to indicate such system events as the mouse pointer crossing a window boundary, and employed redundant visual information to assist the user in locating the pointer when it became lost.

Their additions were well received in user evaluations. Another example not designed specifically for low vision users is Gaver's (1989) Sonic Finder, which provided a range of auditory feedback to users interacting with a desktop environment. Similarly, "force-feedback" mice are available that provide tactile information when the mouse pointer passes into or out of a screen region such as a window, button, or menu item (e.g. Logitech 1998).

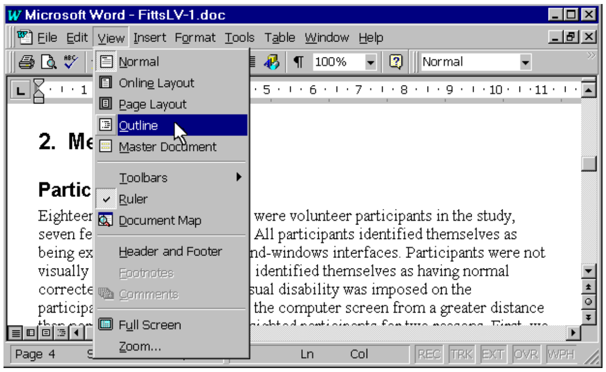

Figure 2. Highlighting of current menu item

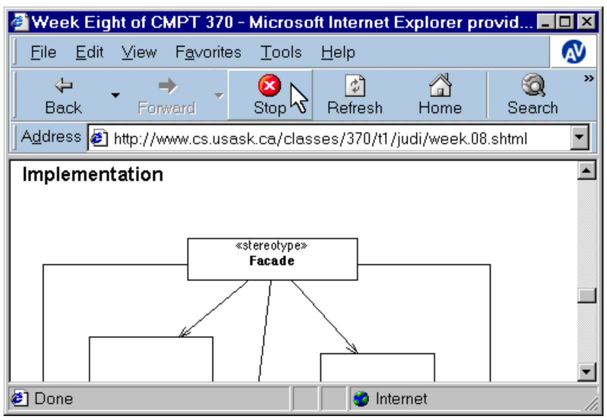

Figure 3. Highlighting (border and colour) of current button on toolbar

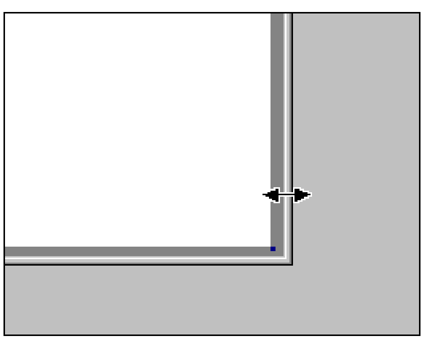

Figure 4. Cursor change over selectable window border

The usefulness of targeting feedback has been studied in a laboratory setting by Akamatsu, MacKenzie, and Hasbrouq (1995), where they examined the effects of different types of targeting feedback on traditional Fitts'- law target selection tasks. The authors looked at four types of feedback: visual, auditory, tactile, and a combination of all three together. Although they did not find a significant improvement in target acquisition times or error rates for any of the feedback conditions, users expressed a preference for feedback over no feedback.

The research carried out by Akamatsu, MacKenzie, and Hasbrouq suggests that redundant targeting feedback is not useful for improving performance. However, their experiments involved normally-sighted users in an optimal viewing environment; we believe that targeting feedback will be of much more value to individuals with reduced visual acuity or individuals operating in a poor visual environment. Our overall hypothesis is that the addition of redundant feedback will make targeting tasks less visually demanding and thus easier for users with limited eyesight. The following sections describe an experiment we carried out to test this hypothesis.

2 Methodology

2.1 Participants

Eighteen undergraduate students were volunteer participants in the study, seven females and eleven males. All participants identified themselves as being experienced with mouse-and-windows interfaces. Participants were not visually disabled; all participants identified themselves as having normal corrected vision. An artificial visual disability was imposed on the participants by having them view the computer screen from a greater distance than normal. We used normally-sighted participants for two reasons. First, it would have been difficult to recruit enough visually-impaired users for the study. Second, visually-impaired users have a wide range of particular visual problems, and the wide variance in the participant population would substantially reduce the precision of our measurements. The ramifications of our choice of participants are discussed in greater detail in later sections.

2.2 Simulated visual disability

Participants were seated further from the computer screen than normal, in order to induce a simulated visual impairment. As discussed above, reduced visual acuity is one major component of visual disability, and visual acuity decreases with distance. Viewing a computer screen from a distance results in many of the types of problems experienced by real low vision users. In the study, participants were positioned by moving them back from the screen until they could no longer read text in the title bar of the application. At this distance, all participants could still see and differentiate between the objects in the test software (start region, target, and mouse pointer). This method of positioning meant that each participant sat at a different distance from the screen; however, most people were placed between five and ten feet. Although the method is imprecise, it does provide a roughly equivalent visual disability for all participants.

2.3 Apparatus

Custom software (see Figure 5) was built for the study to allow participants to carry out standard two-dimensional Fitts'-law tasks. The software was implemented with Tcl/Tk (Ousterhout 1995) and the SNACK sound extension (Sjölander 1999). The software displayed two circles on a grey background: a start region and a target. The start region was 40 pixels in diameter; the target was six pixels in diameter. Target size was chosen to roughly match the size of the smaller targets in standard Windows applications (see discussion above). The display was a nineteen-inch monitor set to a resolution of 1024 x 768 pixels; this means that the targets were 2.1 mm in actual diameter.

A single trial consisted of moving the pointer to the start region and clicking the mouse button, then moving the pointer to the target and clicking again. If the target was successfully selected, the start region was redrawn at the pointer's current position, the target would be redrawn at a new location, and a new trial would begin. Participants were instructed to continue attempting to select the target until they were successful.

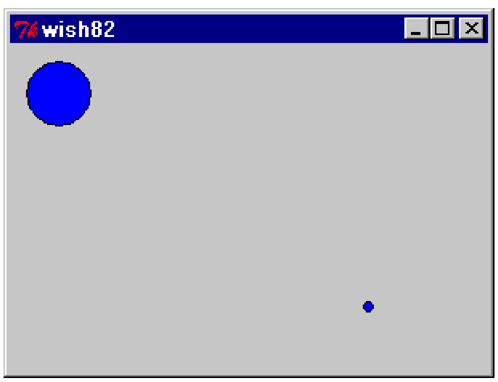

Figure 5. Target-selection software used for the study. The large circle is the start region; the small circle is the target (the window is considerably reduced in extent).

2.4 Study design and conditions

Three experimental conditions were implemented, providing different types of targeting feedback:

- No additional feedback: the target did not change when the mouse pointer entered it.
- Visual feedback: the target changed from blue to red and was highlighted with a red circle whenever the mouse pointer was inside the target (see Figure 6).
- Auditory feedback: when the pointer entered the target a 440Hz tone (approximately "tock") was played for a duration of 0.009 seconds; on exit, a 1760Hz tone (approximately "tick") was played for the same duration.

We used a repeated-measures within-participants design, where participants carried out trials in each of the three experimental conditions. Condition order was counterbalanced. We had three hypotheses in the study:

1. targeting feedback will reduce completion time in targeting tasks
2. targeting feedback will reduce errors (incorrect selection) in targeting tasks
3. targeting feedback will be preferred by participants over no feedback

However, we did not have a prior hypothesis about which type of feedback (visual or auditory) would be more beneficial.

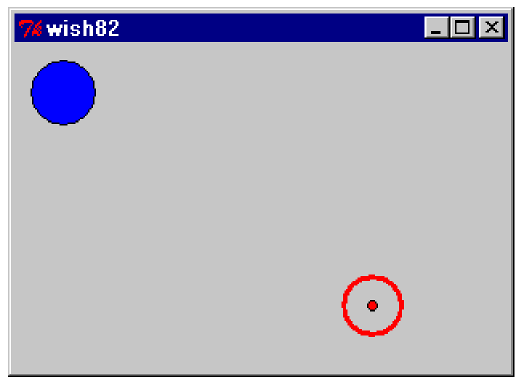

Figure 6. Visual feedback provided when the pointer was inside the target (pointer not shown).

2.5 Procedure

The experimenter positioned the participant at a suitable distance from the screen (as described above) and introduced the task. Participants completed 25 practice trials with no feedback to learn how the trials would work. Participants then carried out trials in each of the three experimental conditions: ten practice trials in order to become accustomed to the feedback, and then 32 test trials in that condition. Participants were instructed to be both as fast and as accurate as possible. The order of the conditions was randomized and counterbalanced. When all the trials were completed, participants were asked questions about their preferences and experiences.

2.6 Data collection

For each trial the following data were recorded: start position and time, end position and time, mouse clicks where the target was missed, and coordinates and timestamps for all mouse moves. After all trials, participants were asked to rank the three conditions in terms of how easy it was to select the targets, and were asked to provide general opinions and comments about their experiences.

3 Results

We collected data to explore each of our three hypotheses—that feedback would improve completion time, error rates, and preference. Our results are organized below into these three areas.

3.1 Completion time

Each participant carried out 32 trials in each of the three conditions. Completion times were calculated using raw start and end times recorded by the software. Times for each set of 32 trials were added together to give a total time for each condition. These data are summarized in Table 1 below and illustrated in Figure 7 (error bars show standard deviation).

Condition	N	Mean	SD
No feedback	18	89.93	21.21
Auditory feedback	18	78.06	12.93
Visual feedback	18	81.91	13.46

Table 1. Mean completion times to carry out 32 trials (in seconds). N = number of participants.

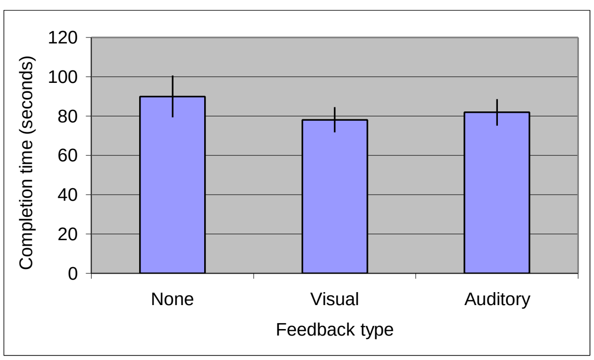

Figure 7. Mean completion times

These data show that participants selected targets more quickly with either visual or auditory feedback. When participants received visual feedback, they completed each trial about three tenths of a second faster than with no feedback; with auditory feedback, they were approximately four tenths of a second faster than with no feedback. Analysis of variance (ANOVA) indicated a main effect of feedback ($F = 5.94$, $p < 0.05$). We conducted followup t-tests on each pair, using a Bonferroni correction to maintain alpha at 0.05. These tests showed significant differences between auditory feedback and no feedback ($p < 0.0167$) and between visual feedback and no feedback ($p < 0.0167$); no difference was found between auditory and visual feedback ($p = 0.026$).

3.2 Errors

Errors were calculated by counting the number of incorrect target selections (mouse clicks). Error data are shown below: Table 2 shows the mean error rate per trial, and these means are illustrated in Figure 8. In Figure 8, error bars show standard deviation.

Feedback type	N	Mean	SD
None	18	10.55	8.55
Auditory	18	3.00	2.30
Visual	18	4.12	3.70

Table 2. Error rates (total errors per 32 trials)

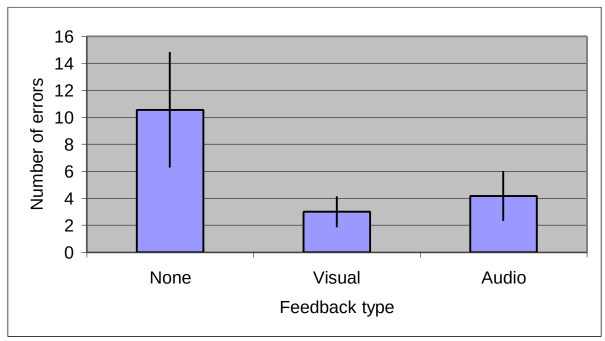

Figure 8. Mean error rates over 32 trials

Participants made fewer errors in selecting the targets when they received visual or auditory feedback. ANOVA again indicated a main effect of feedback ($F = 11.72$, $p < 0.05$). Followup t-tests showed significant differences between auditory feedback and no feedback ($p < 0.0167$) and between visual feedback and no feedback ($p < 0.0167$); no difference was found between auditory and visual feedback ($p = 0.11$).

3.3 Preference

Preferences were determined by asking participants to rank the three conditions in order of preference once all trials were complete. Table 3 below shows the number of participants who placed the different conditions as first, second, and third in their rankings. Figure 9 shows the totals for participants' top preference.

As is obvious from the data, the number of participants preferring some type of feedback is significantly larger than those preferring no feedback. In addition to these results, informal observations of frustration (e.g.

swearing, exclamations) were much more frequent when participants had no targeting feedback.

Feedback Type	First choice	Second choice	Third choice
None	0	3	15
Auditory	6	9	3
Visual	12	6	0

Table 3. Participant preference (cells show number of participants)

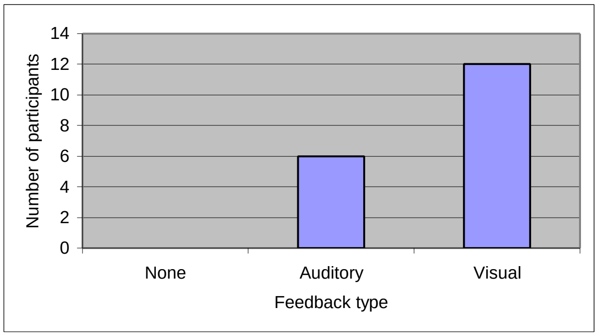

Figure 9. First preference (number of participants)

4 Discussion

Our overall hypothesis was that targeting feedback would assist users who were in a visually stressed condition with targeting tasks. Our results indicate that when users are positioned at a distance from the screen, both visual and auditory feedback lead to decreased task completion time, decreased errors, and increased satisfaction. Targeting time was reduced by three to four tenths of a second for a single task, and errors were reduced from one error in every three targeting attempts to about one in every ten. These results stand in contrast to those of Akamatsu, MacKenzie, and Hasbrouq (1995), where no significant benefits of feedback were found under optimal viewing conditions. In the following sections we discuss reasons why feedback was found to be useful in this situation, comment on how the results will generalize to real users, and suggest guidelines for adding targeting feedback to real-world interfaces.

4.1 Explanation of the results

In general terms, feedback works by providing an indication that some particular event has occurred. In this study, that event was the correct positioning of the mouse pointer over the target. The extra feedback provided in the auditory and visual conditions could lead to improved performance only when the normal feedback—that is, the visual images of the pointer and the target—was inadequate for determining that the pointer was correctly positioned. This was the case for the participants in the study; it was clear from our observations that they found it more difficult to see the pointer and the target from a distance.

The most likely reason why we found an effect when previous studies did not is that previous targeting tasks did not approach this threshold of difficulty. People will generally maintain performance in a test situation simply by expending more effort on the task (Monk 1995); however, it appears that our tasks were difficult enough that additional effort alone could not maintain performance without additional feedback. All participants in our study sat at least twice as far from the screen as they normally would, and so they were working with objects that were considerably reduced in actual size. Doubling the viewing distance reduces the angular width of an object to half of normal and the area of the object to one quarter of normal. In addition, this reduction applies equally to the target and to the mouse pointer, further complicating the task.

Although people do not often sit far away from their displays in real life, our results suggest a continuum of difficulty where redundant feedback can be valuable to users carrying out demanding targeting tasks. The next questions to be answered are whether these kinds of demanding tasks will ever occur in the real world, and if so, what kind of feedback should be provided.

4.2 Generalizing the results to the real world

We consider the generalization of our results to two communities: first, low-vision users, and second, normally-sighted users. In each discussion, we consider how often users will have difficulty with targeting tasks, and consider how well the study participants correspond to users in the real world.

Low-vision users

Reduced visual acuity is one of the defining characteristics of a low-vision computer user. All users with moderate or severe vision loss (20/70 to 20/400) will experience difficulty in determining whether an onscreen pointer is correctly positioned over a target; therefore, targeting feedback should be particularly useful for this group and should increase in value as acuity decreases. For many low-vision users, it is *always* difficult to determine target acquisition with ordinary interface objects, and so any strategy that lessens perceptual demands will be extremely useful.

Although the simulated visual impairment used in this study does not perfectly capture the experience of users with actual impairments, we believe that our results are applicable to the low-vision user community. We chose to manipulate visual acuity specifically because it is defined in terms of distance and therefore affects all users regardless of their eyesight. Reduced acuity is a component of many visual disabilities, and so solutions that aid acuity-sensitive tasks like targeting are worth further attention, even if they must be considered in conjunction with other demands presented by a specific visual condition. As an aside, we note that many of the normally-sighted participants in the study began to exhibit behaviour and characteristics common to low-vision users (e.g. Fraser 1998) when seated at a distance from the screen. For example, people would trap the mouse pointer in the corner of the screen when it became lost, or would shake the mouse to locate the cursor. In addition, participants commented that their eyes became tired quickly and that it was hard to see the white mouse pointer against a grey background; several people stated that they would have preferred a darker background, a scheme used widely by low vision users.

Normally-sighted users

Our results can also be used to inform the design of interfaces for normally-sighted users. Although the particular situation used in our study—users sitting at a distance from the computer—will rarely happen in real life, normally-sighted users often impose other kinds of artificial visual impairments on themselves by operating computers in less-than-optimal visual settings.

As Newell (1995) states, ordinary users often find themselves in situations where they are artificially disabled by factors in the environment. A variety of these factors can contribute to reductions in visual acuity, including fatigue, eyestrain, incorrect eyeglass prescription, fuzzy or flickering displays, poor contrast, or screen glare from overhead lighting. These environmental conditions put people into a visually stressed condition similar to that imposed on the participants in this study, reducing their abilities to see mouse pointers and small targets. Any computer user who has stared at a screen all day after a poor night's sleep can attest that the mouse cursor is not always as visible as it should be. In these situations, our experimental results are likely to generalize well.

Artificial visual impairments, however, are unlikely to profoundly affect acuity; therefore, normally-sighted users are most likely to see benefits in targeting feedback for small targets that are closer to the limit of the user's visual abilities. Our discussions with users suggest that targets of about the size we studied (2-3 mm wide) are good candidates for targeting feedback: this includes interface elements such as object handles, window borders, small icons, and ruler markers. Of these elements, only window borders currently provide feedback in the Windows environment. In addition to small targets, feedback becomes more useful when the cost of error—that is, the cost of erroneously clicking outside the target—is high. For example, there is a considerable cost in mistakenly clicking the close button of a Windows application instead of the maximize button (see Figure 1); in these cases, feedback can improve usability even with larger targets.

4.3 Type, amount, and presentation of feedback

There are many ways to provide targeting feedback, varying widely in type, amount, and presentation. Our discussions with the study participants suggest that flexibility and subtlety will be extremely important considerations when adding targeting feedback to real world applications.

A majority of the participants preferred visual feedback to auditory feedback, but many people in both camps had extremely strong preferences that are not represented in our quantitative data. Some participants were adamant that auditory feedback was the easiest condition to work under, while others remarked that the sound was distracting and made it difficult to concentrate. One user remarked that the visual feedback was too dramatic, and also somewhat distracting. These strong views suggest that feedback type should be controllable by the user. Most people were equally discerning in discussing the amount of information that should be part of targeting feedback (e.g. duration of a sound, area of a visual signal) and its presentation characteristics (e.g. volume, pitch of a sound; colour of visual feedback). There was general agreement, however, that feedback should be subtle, for the simple reason that the user's attention is already focused on the act of acquiring the target.

Since all of our participants were normally-sighted users, we cannot comment on the characteristics of feedback that will be most successful for low-vision users. This question is part of our future work in the area; however, given the wide variety of visual disabilities, it is unlikely that a single inflexible solution will be appropriate for the low-vision community.

5 Future work

Our next steps in this area will be to test the idea of redundant targeting feedback in realistic applications and with real low vision users. Realistic applications will likely demand a more subtle approach to targeting feedback, and we will re-test our hypotheses under these conditions. We are currently building a simple word processor in which targeting feedback will be available for a variety of interface elements.

One additional question to be considered in this application is how to present targeting feedback when there are multiple potential targets (such as the buttons on a tool palette). To avoid confusing or distracting the user, the system must determine when a targeting action is taking place, and only provide feedback for the intended target. Pointer velocity is one possible indicator of the stage of targeting. When complete, the word processor application will be used in a longer-term realistic evaluation of targeting support, with the participation of real low-vision computer users. Although the work described here focuses on support for the acquisition phase of targeting, we are also planning to include support for locating and moving the mouse cursor.

6 Conclusion

In this paper we considered the problem of targeting, where interface elements are selected with an on-screen pointer. Under visually stressed conditions caused by small targets, visual disability, or environmental factors, targeting can become a difficult task. We hypothesized that targeting feedback could improve performance in these conditions. In contrast to studies using optimal viewing conditions, we found that auditory and visual feedback led to improved performance time and lower error rates. We conclude that targeting feedback has considerable potential both for low-vision computer users and for the normally-sighted community as well.

7 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, 816-827. 1995.
2. Baecker, R., Buxton, B., Grudin, J., and Greenberg, S. Chapter Introduction: Touch, Gesture, and Marking. In: *Readings in human-computer interaction (2nd ed.).* R. M. Baecker, W. A.S. Buxton, J. Grudin, and S. Greenberg eds. Kaufmann, Los Altos CA, 1995.
3. Fraser, J.D. *Human Computer Interaction and the Low Vision User.* Unpublished Report, Technical University of Nova Scotia, Halifax NS, 1998.
4. Gaver, W.W. The SonicFinder, An Interface that Uses Auditory Icons, *Human Computer Interaction*, 4, 67-94, 1989.
5. Jacko, J.A, and Sears, A. Designing Interfaces for an Overlooked User Group, *Assets 98: Preceedings of the Third International ACM Conference on Assistive Technologies*, (Marina del Rey CA, April 1998), ACM Press, 75 - 77.
6. Kline, R.L., and Glinert, E.P. Improving GUI Accessibility for People with Low Vision, *Proceedings of ACM CHI'95 Conference on Human Factors in Computing Systems* (Denver CO, May 1995), ACM Press, 14 - 121.
7. Levack, N. *Low Vision : A Resource Guide with Adaptations for Students with Visual Impairments*. Texas School for the Blind and Visually Impaired, Austin TX, 1994.
8. Logitech Corporation, Logitech WingMan Force-Feedback Mouse. Press release at: http://www.logitech.ch/de/about/al_006_29.html
9. MacKenzie, I.. S. Movement Time Prediction in Human-Computer Interfaces. In: *Readings in human-computer interaction (2nd ed.).* R. M. Baecker, W. A.S. Buxton, J. Grudin, and S. Greenberg eds. 483-493. Kaufmann, Los Altos CA, 1995.
10. Monk, A., McCarthy, J., Watts, L., and Daly-Jones, O., Measures of Process, in *CSCW Requirements and Evaluation*, P. Thomas ed., 125-139, Springer-Verlag, London, 1996.
11. Newell, Alan F. Extra-Ordinary Human Computer Interaction, in *Extra-Ordinary Human Computer Interaction*, A.D. Edward ed., Cambridge University Press, New York NY, 1995.
12. Ousterhout. Tcl and the Tk Toolkit. Addison-Wesley Publishing Co., Reading MA, 1995.
13. Sjölander, Kåre. The Snack Sound Extension for Tcl/Tk 1999. www.speech.kth.se/SNACK/
14. Vanderheiden, G. Nonvisual Alternative Display Techniques for Output from Graphics-Based Computers. Journal of Visual Impairment and Blindness, 1989.