The Effects of Feedback on Targeting with Multiple Moving Targets

David Mould and Carl Gutwin

Department of Computer Science, University of Saskatchewan
57 Campus Drive, Saskatoon, SK, S7N 5A9
{mould | gutwin}@cs.usask.ca

Abstract

A number of task settings involve selection of objects from dynamic visual environments with multiple moving targets. Target selection is difficult in these settings because objects move, because there are a number of distracter objects for any targeting action, and because objects can occlude the target. Target feedback has been suggested as a way to assist targeting in visual environments. We carried out an experiment to test the effects of visual target feedback. We found that targeting does become more difficult as the number and speed of objects increases, and that feedback can improve error rates. When feedback was provided on all objects in the space, performance improved significantly over no feedback. Target-only feedback, however, was not significantly better than no feedback. This is a valuable result because all-object feedback is in most cases the only implementation option – since it is usually not possible to pre-determine the user's target among the set of objects.

Key words: Targeting, dynamic environments, moving targets, occlusion, target feedback, Fitts' Law.

1 Introduction

A variety of human-computer tasks involve either (or both) moving targets and target occlusion. Games, especially real-time strategy games, often require the user to select individual units from a crowded battlefield. Scientific visualization or 3-D renderings of geometric data can produce complex pictures from which a user may wish to select a single data point or object. In these settings, targeting can be difficult.

In some cases, objects move autonomously (e.g., command and control situations such as air traffic control or strategy games, as in Figure 1); in other cases, the objects 'move' only through the changing perspective of the viewer (e.g., virtual reality environments or 3D visualizations, as in Figure 2). In both settings, however, object manipulation is usually carried out by selecting objects from the screen with a 2D pointing device such as a mouse. Difficulties in targeting arise because the target moves, and because other objects can occlude the target by moving in front of it.

Figure 1. A screenshot from the real-time strategy game Warcraft3 (www.blizzard.com), with a number of units moving about and potentially occluding one another. Flying units in particular can occlude others.

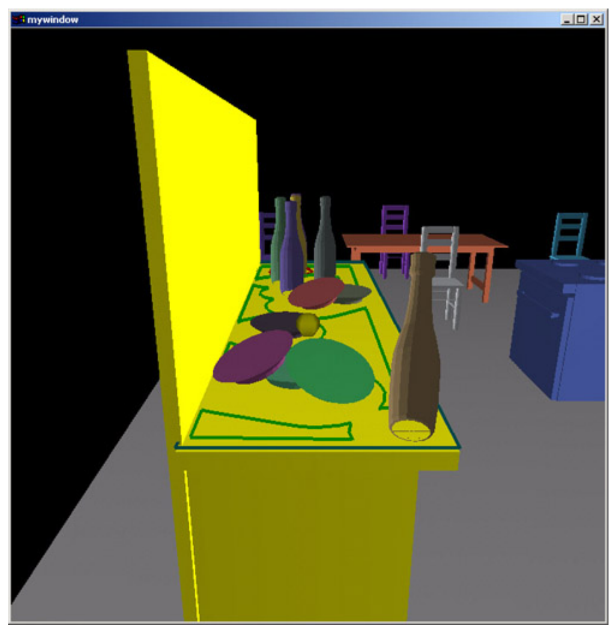

Figure 2. A scene in a VR System (from [13]). As the camera moves, the visual representations of objects move relative to one another.

Feedback has been suggested as a potential way to improve targeting performance [1,2,4]. In a dynamic visual environment, feedback might help users to determine when they are in fact overtop of their intended target. We hypothesized that the utility of feedback would increase as task difficulty grew: while comparatively easy tasks have in the past shown no improvement, more difficult tasks might benefit from feedback. We carried out a study in which people selected a moving target from amongst a collection of other objects, and tested the effects of visual feedback (highlighting the selected object). Both the target and the occluders moved according to predictable paths; the target, in a circular arc, and the occluders, in straight lines. We had three feedback conditions – no feedback, feedback on the target only, and feedback on all objects.

We found that error rates decreased significantly with feedback. Surprisingly, feedback on all objects was significantly better than either of the other two conditions, and target-only feedback was only better than no feedback when there were a large number of other objects on the screen. Our results show that in dynamic environments, having multiple sources of feedback other than the intended target is not a hindrance, and in fact allows for more accurate targeting performance.

In the following sections, we review background research on targeting, moving targets, and target feedback, and then report on the targeting study and discuss implications of our findings.

2 Background

2.1 Targeting actions

Targeting is the act of pointing to and selecting an object on the screen [9]. The stages of a targeting action arise from human movement mechanics, which for targeting can be described by the *optimized initial impulse* model [11]. McGuffin [10] describes the model:

> …an initial movement is made towards the target. If this movement hits the target, then the task is complete. If, however, it lands outside the target, another movement is necessary. This process continues until the target is reached. …In essence, this means that most aimed movements consist of an initial large and fast movement that gets the subject reasonably close to the target, followed by one or more shorter, and slower, corrective movements that are under closed-loop feedback control. (p.18)

When the pointing device in the interface has an on-screen pointer (as opposed to a touch screen), we can thus divide targeting into three distinct stages: locating, moving, and acquiring [4]. *Locating* is the initial act of finding the pointer on the screen when its position is unknown. *Moving* is the act of bringing the pointer to the general vicinity of the target, requiring that the user stay aware of the pointer's position as it travels across the screen. *Acquiring* is the act of precisely setting the pointer over the target and determining that the pointer is correctly positioned.

2.2 Target feedback

Although graphical interfaces already provide targeting feedback (by showing the cursor moving across the screen and onto the target), additional feedback has been considered in several circumstances. The additional feedback is usually aimed at the acquisition stage of targeting, and is most often visual, although auditory and tactile feedback have also been considered.

A common visual technique involves highlighting an 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 is used in menus and toolbars where individual items do not have a clear visual boundary. Auditory and tactile targeting feedback has been used in research prototypes for low-vision users (e.g., [7]), and in some commercial devices, such as force-feedback mice that bump or vibrate as the cursor moves over object boundaries.

However, the value of feedback on targeting performance has not been clearly established. Previous studies using normal targeting conditions [2] have found no effect of visual, auditory, or tactile feedback (or of a combination of all three), although participants preferred the feedback condition. These studies, however, were done with single unmoving targets, and were carried out under normal visual conditions.

One previous study that did find an effect of targeting feedback involved a simulation of low-vision conditions [4]. Participants were seated at a distance from the computer screen to artificially reduce visual acuity. In this situation, where targets and pointers were difficult to see, both auditory and visual feedback produced significant differences in targeting performance compared with no feedback.

These previous investigations suggest that in most cases, the ordinary visual feedback provided by the interface is adequate for maximizing targeting performance. Adding additional feedback in these cases will have little if any effect since performance is already nearly optimal. However, in situations where optimal performance is difficult, additional feedback has the potential to improve performance. It is possible that targeting in dynamic environments, with moving targets and multiple distracters, is of this type.

One additional issue also raised by these previous studies is the potential for 'feedback confusion' with multiple targets. When there are several objects that could be selected by the mouse, all of them would have to provide targeting feedback, potentially causing distraction [2] or even the loss of any benefit for selecting the intended target [4]. Although this does not appear to cause problems in menus, there is perhaps a higher likelihood of distraction when targets are more densely clustered in a 2D space.

2.3 Moving targets

Several previous researchers have considered the factors involved in targeting moving objects, although there are few studies that specifically involve human-computer tasks with 2D interfaces. Much of the work has been done in real-world settings where tasks with moving targets (e.g., catching a ball) are common.

There are reformulations of Fitts' law that involve targets moving at constant velocity (e.g., [5,6]). These models incorporate the target's velocity into a new index of difficulty for targeting tasks, where higher velocities imply more difficulty in targeting. However, these models generally consider only one-dimensional movement (towards or away from the cursor) [5].

Intercepting a moving object involves different mechanics and additional factors compared to static target selection. According to Fitts' law, targeting difficulty for unmoving objects is determined completely by the size of the target and its distance from the starting location (e.g., [9]). The key difference with moving targets is that the timing of the acquisition action becomes a factor [3,8]. With static targets, only the position of the target must be considered in planning (or adjusting) a targeting motion. With moving targets, however, both the position of the target and the time that the cursor arrives at that position are important. A targeting motion could be 'correct' in terms of distance and position, but if the mouse button is pressed at the wrong time, the target will either not have reached the cursor, or will have already passed by.

There is a tradeoff between speed of motion and timing, just as there is a tradeoff between speed and accuracy of positioning (as in static-target models such as Fitts' law) [3]. This tradeoff is optimized differently for different target speeds; researchers have noted that people use different hand speeds to intercept targets that are moving at different speeds [3].

An alternate strategy for selecting moving targets is to move to a point on the target's trajectory, and lie in wait for it (i.e., ambush the target rather than chase it). In these cases, the target moves through an 'interception zone' [12] and the user's task becomes entirely one of timing rather than positioning. Models have also been developed that can predict performance in this version of the task.

To further explore the effects of visual target feedback in with moving targets, and in environments where there are multiple objects, we carried out the experiment described below. We examined the effects of three factors: the number of objects on the screen, the speed at which the objects move, and the presence of visual targeting feedback.

3 Study Methodology

The following sections provide details about the study participants, the apparatus and tasks used, the experimental factors, and the study design.

3.1 Participants and apparatus

Eighteen people (14 male, 4 female) were recruited from the computer science department of a local university, and were given course credit for participating in the study. All participants were frequent users of mouse-and-windows based systems (at least 20 hours per week). Seven participants were gamers who played at least one hour per week of real-time strategy, first-person shooter, or role-playing games.

The experiment was conducted on a 2.8 GHz P4 Windows XP PC running a custom OpenGL application (see Figure 3). The display was a 15" monitor set to 1024x768 resolution; the window for the study system was 600x600 pixels, and was centred on the screen.

3.2 Tasks and experimental conditions

The system presented two-dimensional target-selection tasks in several different occlusion and movement conditions (see Figure 3). At the beginning of each trial, participants positioned their mouse cursor in a blue square at the lower left of the screen. The system then displayed a moving target, coloured white and marked with a green 'X'. To ensure that participants saw the target when it appeared, the lines of the target's X were initially extended 100 pixels beyond the borders of the target (extensions were removed after three seconds). Further, the color of the home square was changed to red at the moment the trial began.

The target moved in an arc around the home square at lower left (i.e., moving from upper left to lower right and back again). In each experimental condition, a set of specific radii were presented in random order. In some conditions, occluder objects could also be present on the screen, and the target was always drawn behind these objects. The participant's task was to move the cursor onto the target and click the mouse button, as quickly and accurately as possible. Only clicks where the cursor was over a non-occluded area of the target were considered to be successful selections; all other clicks were considered to be errors.

The study involved three factors: feedback type, number of occluders, and object velocity.

- *Feedback*. Our main research question considered the effects of targeting feedback. Feedback in the study involved changing the colour of an object (to bright green) when the mouse cursor was positioned above a visible part of that object. Two types of feedback were possible (in addition to no feedback at all). First, feedback could be given when the cursor was over the target only. Second, feedback could be given when the cursor was over any screen object (both target and occluder objects).

- *Number of occluders*. To determine whether feedback effects vary by the number of other objects that could occlude the target, we tested a range from few (22) occluders to a medium amount (44) to many (88). Examples are shown in Figure 4.

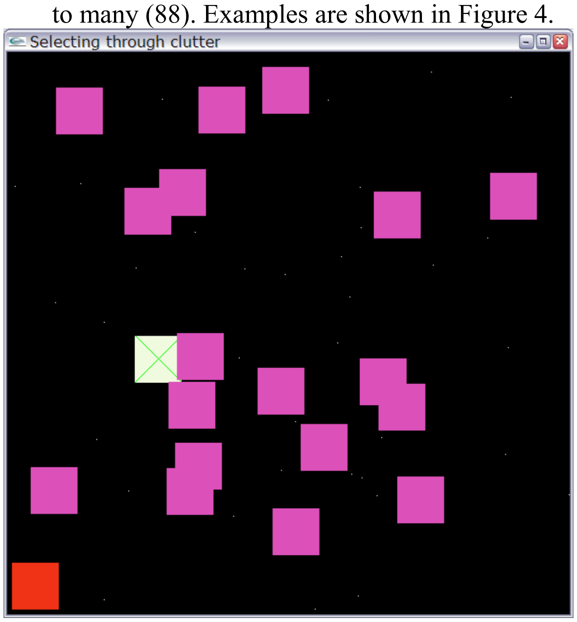

Figure 3. Example screen from study system. Box at lower left is the start area; target is marked with an X.

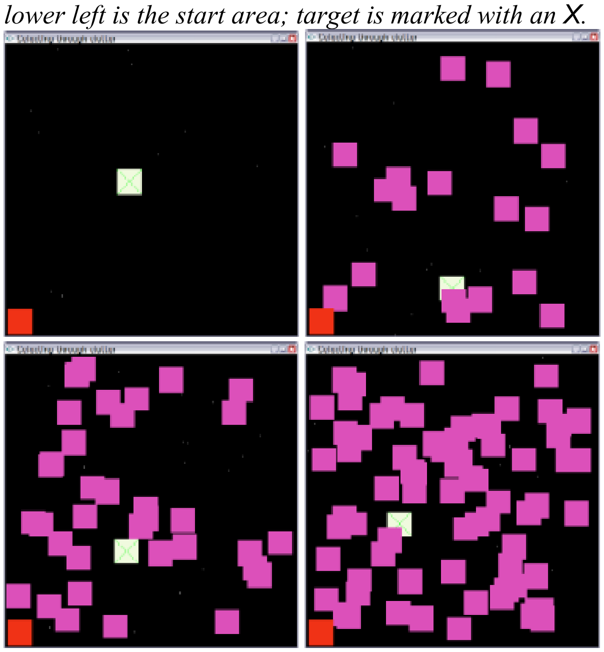

Figure 4. Number of occluders: 0, 22, 44, 88.

- *Velocity of objects*. Since occlusion of the target occurs because the other objects are moving, we tested three different movement speeds: slow (45 pixels/sec), medium (220 pixels/sec), and fast (400 pixels/sec). Figure 5 shows relative distances for these speeds.

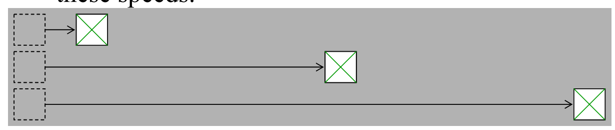

Figure 5. Velocity conditions. The diagram shows the relative distance that objects move in one second: 45 pixels/sec (top), 220 pixels/sec (middle), 400 pixels/sec (bottom).

3.3 Experimental design

The study used a 3 x 4 x 3 within-participants factorial design. The factors were:

- Feedback type: none, all objects, target only
- Number of occluders: 0, 22, 44, 88
- Object velocity: 45, 220, 400 pixels/sec.

Feedback type was fully counterbalanced; the other two factors were always presented in increasing order (i.e., from fewer to more occluders, and from slower to faster). Within each condition, participants carried out 16 targeting trials.

With 18 participants (3 in each order group) and 16 trials per condition, there were 10368 trials recorded. The study system collected completion times and error information for each target, and recorded the amount of the target that was visible when the participant selected it. In addition, answers to summary questions were recorded on a questionnaire.

3.4 Procedure

Participants were assigned to one of six groups to determine the order of the feedback conditions. Participants were then introduced to the experiment and to the study system, and were asked to complete four practice trials in each condition. Participants then completed 16 targeting tasks in each of the 36 study conditions. Participants were instructed to click on the targets as quickly and as accurately as possible. Rests were allowed between conditions. After all conditions for a session were complete, participants were asked three questions: the type of feedback with which they believed they were fastest; with which they believed they were most accurate, and which they preferred overall.

4 Results

The following sections report our findings for the three primary factors in the study (feedback, number, and velocity) in terms of our three dependent variables (time, errors, and amount of occlusion at selection time)[1]. Findings from the preferences survey and from an analysis of how gaming experience affected performance follow the main results.

4.1 Effects of feedback, number, and velocity

Errors

Using 3x4x3 Analysis of Variance (ANOVA), we found significant main effects of each factor on error rate: for feedback type, $F_{2,34}=10.75, p<0.001$; for number of occluders, $F_{3,51}=23.62, p<0.001$; for velocity, $F_{2,34}=49.40, p<0.001$. Means for each factor are shown in Figures 6-8.

We used Tukey HSD tests to look for differences between the different levels of each factor. As expected, there were significant differences between the different levels of both object velocity and number of objects. More surprising were the differences between the feedback conditions. There were significantly fewer errors with all-object feedback than with either of the other types (for no feedback, $p<0.001$; for target-only feedback, $p<0.01$). However, there was no difference in error rate between target-only feedback and no feedback ($p=0.61$). The differences mean a reduction of approximately eleven percent for all-object feedback compared to target-only feedback, and of about 32 percent compared to no feedback.

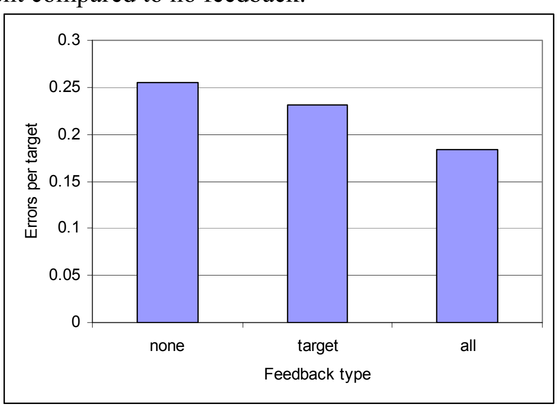

Figure 6. Mean errors, by feedback type.

The only significant interaction between factors was between feedback type and the number of objects ($F_{6,102}=4.07$, $p<0.005$). As can be seen from Figure 8, the differences between feedback and no feedback increases as the number of objects increases. To examine the interaction more closely, t-tests were carried out between pairs of conditions for each number group. Significant differences were found only with 44 and 88 objects. With 44 objects, feedback on all objects was significantly better than both no feedback ($p<0.005$) and target-only feedback ($p<0.001$). With 88 objects on the screen, both of the feedback conditions were better than no feedback (for all feedback, $p<0.001$; for target-only feedback, $p<0.01$). However, no differences were found between the two feedback conditions with 88 objects ($p=0.22$).

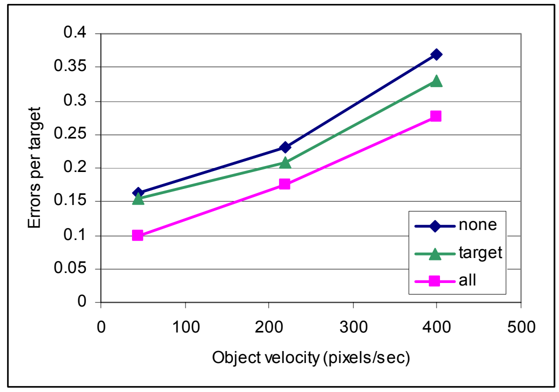

Figure 7. Mean errors, by object velocity and feedback type.

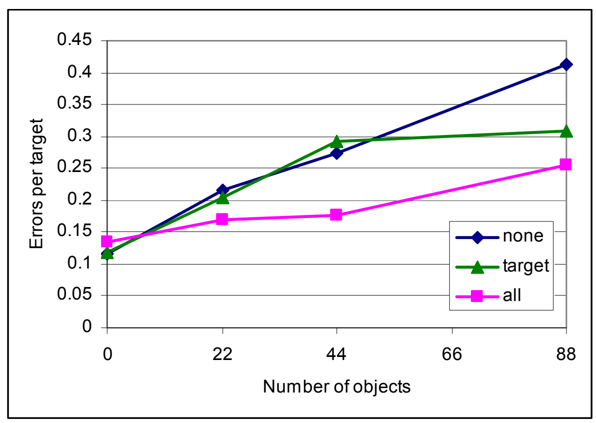

Figure 8. Mean errors by number of objects and feedback type.

Amount of occlusion

This measure recorded how much of the target was visible when the participant actually acquired it. Our expectations were that with more objects on the screen, average target occlusion would increase (because targets would be more likely to be covered), and that with target feedback, participants would be able to select targets that were more covered than without feedback.

There were main effects of number ($F_{3,51}=1714.45$, $p<0.001$) and object velocity ($F_{2,34}=14.34$, $p<0.001$). However, there was no main effect of feedback type ($F_{2,34}=0.58$, $p=0.565$), implying that participants did not use the targeting feedback as a way to select less-visible targets (see Figure 9).

[1] One participant's data was removed from the study due to his difficulty in following the given instructions.

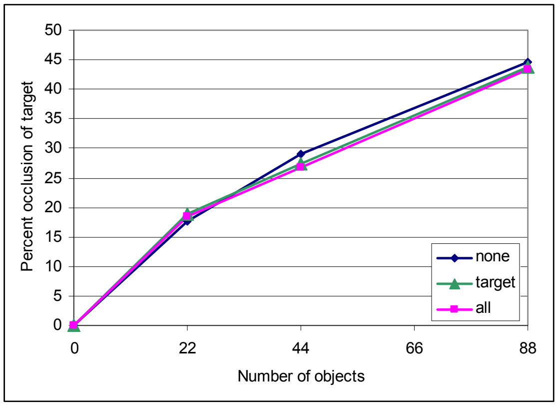

Figure 9. Target occlusion amount, by feedback type and number of objects.

Completion time

Targeting trials took on average about one second. There were main effects of number and velocity on completion time: for number of objects, $F_{3,51}$=210.88, p<0.001; for object velocity, $F_{2,34}$=37.24, p<0.001. Figures 10 and 11 show these data. There was no effect of feedback type on completion time ($F_{2,34}$=0.36,p=0.70), and mean times for the three types were all very close – within 10ms of each other.

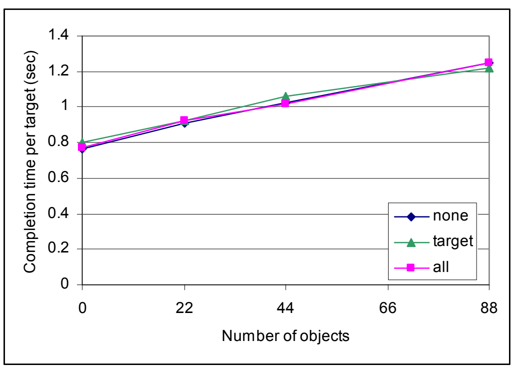

Figure 10. Completion time, by feedback type and number of objects.

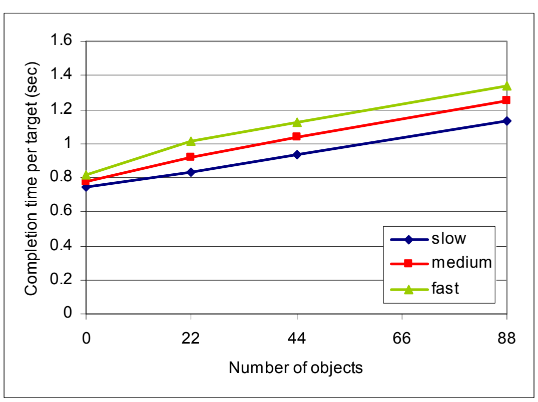

Figure 11. Completion time, by target speed (45, 220, 400 pixels/sec) and number of objects.

4.2 Effects of gaming experience

The performance of the seven participants with game-playing experience was compared with the non-gamers. There was an effect of game experience on completion time ($F_{1,16}$=7.75, p<0.05): regular game players completed trials (on average) in 0.88 seconds compared to 1.08 seconds for non-gamers. However, there were no effects on error rate ($F_{1,16}$=0.42, p=0.53) or target occlusion at selection time ($F_{1,16}$=0.10, p=0.76); there were also no interactions between completion time and the other factors.

4.3 Preferences

Three questions were given to participants at the end of the session, concerning speed, accuracy, and overall preference. People strongly preferred feedback on targets only, as shown in Figure 12.

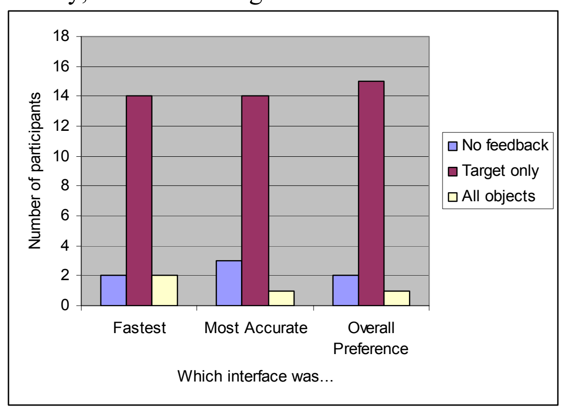

Figure 12. Participant preferences

5 Discussion

Our main hypotheses in the study were: that targeting becomes more difficult in environments with more objects, and where objects move more quickly; and that targeting feedback should make more of a performance difference the more difficult targeting becomes. The study found support for these hypotheses. Our main findings are:

- Completion time and errors increase both with the number of objects on the screen and with the speed of the objects.
- Target feedback had no effect on completion time in any condition.
- Target feedback reduced error rates, and was increasingly helpful as the number of occluders increased.
- Feedback on all objects was significantly better than feedback on the target alone; target-only feedback was only better than no feedback when there were the highest number of objects on screen.

- There was no effect of feedback on the amount of occlusion at the moment of target selection.
- Users strongly preferred target-only feedback to either no feedback or all-object feedback, and believed that this condition was fastest and most accurate.

5.1 Explanations for the results

Completion time

We expected that time needed to complete a task would increase as the difficulty of the task increased. Increasing the speed of the target is known to increase targeting difficulty [6]. It is also reasonable that increasing the number of occluders increases difficulty, both because target acquisition becomes more challenging (locating the target in a cluttered environment is more difficult than in a clean environment) and because the selectable area of the target may be reduced by occlusion. An obscured target is more difficult to hit because it is less visible and because less of the target is available.

We characterized targeting actions as a sequence of locating, tracking, and acquiring stages. The type of feedback we used (object highlighting) mainly helps in the last of these phases; certainly, target-only feedback is effective only in the acquisition phase. However, because feedback on all objects provides an element of tracking feedback (i.e., when there are enough occluders, the highlighting shows the path of the cursor across the screen), we speculated that all-objects feedback might help the user to locate the mouse pointer when the pointer was off-target. By illuminating the occluder that contained the mouse pointer, the user got a strong (albeit erratic) signal about the pointer's position. However, the data showed no effect of feedback type on completion time.

Error rates

We found that the error rate increased as task difficulty increased. We defined an error as any mouse click that happened while the mouse pointer was outside the visible target area. We attribute the increase in error rate with target speed to the propensity of the moving targets to escape the pointer, with the likelihood of escape increasing as the speed increased. Similarly, we attribute the observed increase in error rate with increased occluder count to an increased likelihood that a click will occur in a portion of the target that is not visible.

The ability of feedback to reduce error rate seems clear, in that feedback provides a strong indication of when acquisition will be successful. If the target is lit, then it is safe to click; if the target goes dark, then it is unsafe. However, this benefit (in the case of target-only feedback) appeared to only take effect at higher numbers of occluders, which is compatible with earlier suggestions that performance can only be improved in more difficult targeting situations.

The difference between all-object feedback and target-only feedback, however, is a more difficult result to explain. One possible mechanism is that users may have utilizing the extra feedback as an enhancement to the acquisition warning system. Target-only feedback gives a signal about whether the target is selectable, but all-objects feedback gives an even stronger signal: not only do we have the warning about missing the target, we have an extra warning from the highlighting of an adjacent occluder.

A second possibility is that the tracking information that is provided by all-object feedback gives some advantage to acquisition. If, for example, people begin their acquisition action (i.e., begin pressing the mouse button) while they are still in the tracking stage, then the provision of tracking feedback could potentially have an effect on accuracy. However, this is a hypothesis that must be tested in further research.

Amount of occlusion

That there was no effect of feedback on amount of target occlusion at the time of selection can be anticipated in view of the absence of any effect of feedback on completion time. If the users were able to select smaller target by using feedback, they should have been able to complete their tasks more quickly on average, either by hitting the target before it became too occluded, or hitting it sooner while it was emerging from occlusion. Either of these would have produced shorter completion times.

Preferences

Users strongly preferred the target-only feedback condition. In written comments, six of eighteen subjects remarked that they found highlighting of occluders in the all-feedback condition to be distracting. We suggest that the users sought to concentrate on the moving target, and that although changes at the periphery of their attention were viewed as irrelevant, the users were in fact able to make use of this information, reducing their error rate. Further, the all-feedback condition, while it failed to improve completion time, at least had no negative effect on timing. One possible reason that target-only feedback was preferred is that it gave only positive advice: the target lit up when the user was successful, and nothing happened otherwise, thus emphasizing the user's successes. This is in contrast to all-objects feedback, where both success at acquiring the target and failure to acquire the target were presented as feedback.

5.2 Implications for practitioners

In a real application, we are not able to provide target-only feedback, because until the user clicks, we cannot identify the intended target. However, it is straightforward to provide all-object feedback. There has been some concern about feedback confusion with multiple targets, but in our data such effects did not manifest themselves – on the contrary, we were able to obtain additional benefit in the form of a reduced error rate.

Thus, the all-object feedback condition is a viable one for real applications, with one caveat. We would expect users' experiences in real applications to be similar to subjects' experiences in our study to the extent that the tasks they undertake are similar. We have suggested a few applications in which users might desire to select among moving targets in a cluttered environment: in VR applications, or visualization of 3D data, or in computer games. In all of these situations we would have many objects, potentially occluding one another, with the possibility of movement. In VR or visualization, the difficult task of selecting a moving target could be simplified by pausing the movement. The VR avatar could stop moving for a moment, then continue after the object is selected; the rotating dataset could be halted while the feature is indicated. In games, although it may be possible to pause before undertaking a targeting task, the desire is to be able to select a target without interrupting the action. Particularly in the case of real-time strategy games, selecting a target from among numerous moving targets is a common task.

6 Conclusions and future work

Targeting in environments with multiple moving targets can be difficult. We carried out a study to determine the effects of target feedback under conditions of motion and occlusion. We found that target feedback improves error rates, and surprisingly, that feedback on all objects helps more than feedback on the target alone. Our results suggest that feedback should be implemented in interfaces that involve selection of moving objects.

There are several issues raised by this study that should be considered in future research. First, we plan to investigate whether we can implement a version of all-object feedback that users find less distracting. It is possible that by toning down the visual effect (or switching to another modality), the distracting aspects of the feedback can be reduced. Second, we will further explore the issue of how feedback in the tracking stage of targeting may influence error rates. Targeting feedback has traditionally only focused on acquisition actions, and it may be possible to extend the value of the approach by supporting more of the overall targeting task.

References

[1] Akamatsu, M., & MacKenzie, I. S., Movement Characteristics Using a Mouse with Tactile and Force Feedback. *International Journal of Human-Computer Studies*, 1996, 45, 483-493.

[2] 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*, 1995, 38, 816-827.

[3] Brouwer, A., Brenner, E., and Smeets, J., Hitting Moving Objects: the Dependency of Hand Velocity on the Speed of the Target, *Experimental Brain Research*, 2000, 133, 242–248.

[4] Fraser, J., and Gutwin, C., The Effects of Feedback on Targeting Performance in Visually Stressed Conditions. *Proceedings of Graphics Interface 2000*, Montreal, 203-210, 2000.

[5] Hoffmann, E. R., Capture of Moving Targets: a Modification of Fitts' Law. *Ergonomics*, 1991, 34, 211-220.

[6] Jagacinski, R., Repperger, D., Ward, S., & Moran, M., A Test of Fitts' Law with Moving Targets. *Human Factors*, 1980, 22, 225-233.

[7] Kline, R.L., and Glinert, E.P. Improving GUI Accessibility for People with Low Vision, *Proceedings of ACM CHI'95,* ACM Press, 14-121, 1995.

[8] Lee, D., Port, N., and Georgopoulos, A., Manual Interception of Moving Targets: II. On-Line Control of Overlapping Submovements, *Experimental Brain Research*, 1997, 116, 421–433.

[9] MacKenzie, I. S., Movement time prediction in human-computer interfaces. In R. Baecker, J. Grudin, W. Buxton, & S. Greenberg (Eds.), *Human-Computer Interaction: Towards the Year 2000*, 1995, San Francisco, Morgan-Kaufmann, 483-493.

[10] McGuffin, M., *Fitts' Law and Expanding Targets: An Experimental Study, and Applications to User Interface Design*. Unpublished M.Sc. Thesis, Computer Science, University of Toronto, 2002.

[11] Meyer, D., Abrams, R., Kornblum, S., Wright, C., and Smith, J., Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review*, 95(3), 340-370, 1988.

[12] Port, N., Lee, D., Dassonville, P., and Georgopoulos, A., Manual interception of moving targets: I. Performance and movement initiation. *Experimental Brain Research*, 116(3), 406-420, 1997.

[13] Xu, K., Stewart, J., and Fiume, E., Constraint-Based Automatic Placement for Scene Composition, *Proc. Graphics Interface 2002*, 25-34.