

Human On-line Response to Target Expansion

Shumin Zhai^ξ Stéphane Conversy[♦] Michel Beaudouin-Lafon[♦] Yves Guiard^{*}

[♦] LRI & INRIA Futurs
Université Paris-Sud, Orsay, France
{conversy, mbl}@lri.fr

^ξ IBM Almaden Research Center
San Jose, CA, USA
zhai@us.ibm.com

^{*} Mouvement et Perception
CNRS & U2, Marseille, France
guiard@laps.univ-mrs.fr

ABSTRACT

McGuffin and Balakrishnan (M&B) have recently reported evidence that target expansion during a reaching movement reduces pointing time even if the expansion occurs as late as in the last 10% of the distance to be covered by the cursor. While M&B massed their static and expanding targets in separate blocks of trials, thus making expansion predictable for participants, we replicated their experiment with one new condition in which the target could unpredictably expand, shrink, or stay unchanged. Our results show that target expansion occurring as late as in M&B's experiment enhances pointing performance in the absence of expectation. We discuss these findings in terms of the basic human processes that underlie target-acquisition movements, and we address the implications for user interface design by introducing a revised design for the Mac OS X Dock.

Keywords

Pointing, Fitts' law, human performance, human information processing

INTRODUCTION

Fitts' law has been one of the cornerstones of HCI research. Its contribution to user interface design and evaluation includes principled evaluation of different input devices [4, 8, 10, 11, 17], systematic comparison of two styles of interfaces such as crossing-based vs. pointing-based interaction [2], and optimization of stylus-based virtual keyboards [12, 18]. Fitts' law has also inspired efforts in discovering other types of regularities in HCI tasks such as path steering [1] and novel interaction techniques such as the area cursor [9].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2003, April 5–10, 2003, Ft. Lauderdale, Florida, USA.
Copyright 2003 ACM 1-58113-630-7/03/0004...\$5.00.

According to Fitts' law [5], target acquisition time (or movement time, *MT*) on a computer screen is determined by the logarithm of the ratio between target distance (*D*) and target width (*W*), known as the task's Index of Difficulty (*ID*). For any system made up of a device, a technique, and a user, the *MT* is a linear function of the *ID*. While this regularity is remarkably general, the intercept and the slope of the linear relationship vary from system to system and thus can quantify performance.

From the point of view of interface design, Fitts' law suggests two non-exclusive design recommendations to facilitate performance. First, the distance that the cursor has to cover should be kept reasonably short, or perhaps shortened at the appropriate time [3]. Second, target size should remain reasonably large, or perhaps expanded at the appropriate time. It is the latter suggestion that is discussed in the present paper.

The concept of an expanding target comes up with the problem of real-estate limitation. If we want larger targets, then we will be able to display fewer items, potentially a concern in the case of multiple, small-size windows. This tradeoff makes the target size problem non-trivial. A provocative effect in this respect was recently reported by McGuffin and Balakrishnan [13] (hereafter M&B). They suggested a way of facilitating pointing with no permanent spatial cost by temporarily expanding the target during the end of the approach. They found that performance benefited from target expansion even when the target only began to expand as late as after 90% of the movement toward the target.

M&B's finding is potentially important both practically and theoretically. Practically it suggests the possibility of 'having the cake and eating it too' in terms of space management on computer screens.

Theoretically, it supports a long line of thinking in human motor control. An old theory there is a two-phase theory in rapid reaching movement – ballistic and current control [16], with the first phase being open-loop and the second closed-loop. A more recent theory is that of [14], which postulates an optimized trade-off between the first and second phase. Since only the second phase involves

feedback control, it is plausible that a target that expands just before the second phase of control is begun is as good as a target that has that larger size permanently.

Since M&B reported no kinematic data, the time durations corresponding to 90 and 10% of movement amplitude are unknown. Given that human closed-loop reaction time takes about 100-200 ms [15], it is surprising that one can still take full advantage of the expanded target at so late a stage in movement execution.

One concern about M&B's study is the treatment of errors. When they had missed one target, M&B's participants were allowed to correct the position of their cursor and click again until a hit was recorded. So an error could only result in a lengthening of *MT*. Such a procedure is convenient because it makes it possible to quantify performance with a single dependent measure. It fails, however, to capture one important aspect of pointing in real GUI's, where the target object is often one element of a pack (e.g., a menu bar): in such a case an error will not simply amount to a miss, but rather to a wrong selection, with a net time cost likely to exceed by a large amount that of a simple aiming correction. Thus, it would have seemed preferable to consider separately the effects exerted by target expansion on the speed and the accuracy of aimed movements.

Finally, an experimental design option in M&B complicates its theoretical implication. Because all trials in the expanded target condition were massed together, the participants were allowed to anticipate the expansion event, whose probability was either zero or one. Thus, the time saving might have been due, to unknown extents, to the user's ability to react on-line to enlarged target size, and/or to their ability to visualize the enlarged target from the beginning of the movement.

Another recent study on changing target acquisition was reported by Gutwin [7], which showed a *negative* impact of target change. In Gutwin's experiment, the target—or more precisely the area under the moving cursor—was subject to fisheye expansion. The stronger the distortion, the more time it took to complete a pointing task, and the more errors (See Figure 7 and 8 of [7]). An important difference between Gutwin and M&B, however, was that both target distance and target size changed in Gutwin's study whereas only target size changed in M&B.

Although rarely done in the field of human computer interaction, important research results should be replicated for soundness and sustained development of a research topic. Our current study revisited M&B experiment, but with an extra condition that had expansion, shrinking, and static target randomly mixed so that the participants could not anticipate the final target size before or in the early phase of reaching movements. We also controlled and measured error rate separately as is more conventionally done in Fitts' law studies. Our aim was to pin down the

mechanism of user's reaction to target expansion and its implications to user interface design

METHOD

Task

To obtain data comparable to M&B's, we introduced as few modifications as possible to their methods.




Figure 1. Experimental Task

We used their discrete, one-dimensional pointing task (Figure 1). The participant had to first move the cursor to a home position marked by a rectangle placed at the extreme left of the screen, rest there for 0.5 sec (reduced from 1 sec in M&B), start moving to the target as soon as it appeared on the right, and click it. The time which participants had to minimize was that elapsed between the presentation of the target and its selection with a mouse click. Whether a hit or a miss (signaled by an acoustic beep), the click terminated the movement. The participants were urged to adjust their speed-accuracy strategy so as to make about 4% errors on average in all conditions. As an aid for this adjustment, the program displayed the error rate after each block of trials, together with slow-down or speed-up suggestions.

Experimental conditions

Static (baseline): target size never changed during the movement.

Expanding: targets always expanded by a factor of two when the cursor had covered 90% of target distance (*D*). On the screen the target gradually expanded in 100 ms, but its final size was instantaneously made available to the cursor. This condition allowed the participants to anticipate final target size. This as well as the preceding condition served to replicate M&B's experiment.

Random: Final target size was unpredictable. After 90% of distance coverage, the target would, with an equal probability of 1/3, expand (twice the initial size), shrink (half the initial size), or remain the same—hereafter we will designate this within-block variable as *Target Type*.

Note that recourse to shrinking targets was necessary to prevent a biased expectation of the final target size.

Task ID

We used fewer *D/W* combinations than M&B (8 rather than 13), but the same range of *ID*, from 3.17 to 7.01 bits with equal 1.3-bit increments. Our eight combinations of *D/W* (in screen pixels) were 512/64, 512/25, 512/10, 512/4, 900/113, 900/44, 900/18, and 900/7.

Participants

Twelve adult volunteers, 11 male and one female, aged 28.8 years on average ($SD = 7.3$ years), served in the experiment. They received 10 Euros for their participation.

Procedure

As this experiment involved one more condition than M&B's and we wished to provide our participants with similar practice times, our procedure had to be slightly different. We also needed a sufficient number of trials per block to make sure that target expansion was truly unpredictable in the random condition (a long enough sequence of equally frequent A, B, and C events is needed for the next event to be difficult to guess). Whereas M&B used 65-trial blocks (5 presentations of each of their 13 *D/W* combinations), we used 120 trial-blocks (15 presentations of each of our eight *D/W* combinations). So the random condition of our design offered, for each *D/W* combination, 15 instances over which target evolution could be randomized (five static, five expanding, and five shrinking targets). We generalized the 120-trial structure to the other two conditions. Our blocks of trials being longer than M&B's, we used fewer, two blocks per condition rather than five, leading to 240 trials per condition and an overall total of 720 trials per participant (650 in M&B).

Each participant performed in all three conditions. The order of the three conditions, each ran twice in two consecutive blocks, was counterbalanced according to a 3×3 Latin square, with four participants assigned to each of the three orders.

Apparatus

We used an equipment similar to that described by M&B, a 21-inch 1280×1024 resolution color monitor, a Wacom Intuos 12×18 inch digitizing tablet with a puck handle set at 1/1 control display ratio.

RESULTS

One systematic option below is that we analyze *MT* and error rate in parallel to make sure that every effect observed on speed is not jeopardized by a contrary effect on accuracy. We consider three dependent variables: *movement time (MT)*, the time elapsed from the beginning of cursor motion to the final mouse click, *error rate*, the number of misses relative to the total number of reaching

movements, and *reaction time (RT)*, the time elapsed between target presentation and the beginning of cursor motion.

Repeated measures analyses of variance were performed on these three variables. We used two main ANOVA designs, one ignoring the change of target size in the Random condition (4 levels of *ID* * 3 Conditions * 2 Blocks) and the other, restrained to the Random condition, taking Target Type into account (4 levels of *ID* * 3 Target Types * 2 Blocks). Unless specified otherwise, the *ID* is calculated according to the initial width of the target.

Effect of Practice on Movement Execution

MT improved from Block 1 (mean 895 ms) to Block 2 (871ms) ($F_{1,11} = 7.61, p = .019$), but this effect did not significantly interact with any other. On the other hand, the Block factor had no significant effect on error rate ($F_{1,11} = .95, p = .35$), nor did this factor significantly interact with any other for accuracy. So we will base our calculations below on the data from both blocks.

The M&B Effect: Facilitating Pointing with Predictable Target Expansion

Here we consider the always Expanding and the always Static conditions, our two conditions that replicated those of M&B.

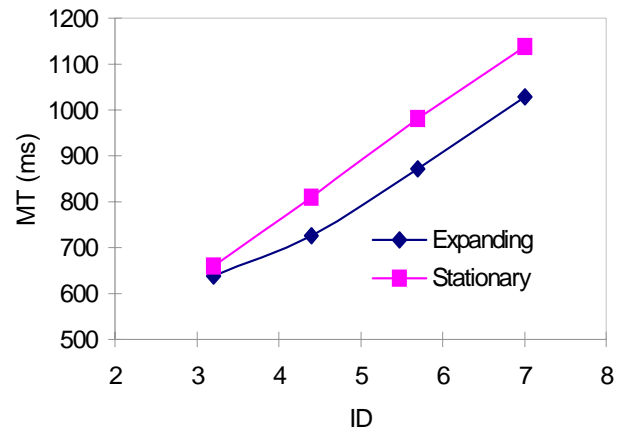


Figure 2. *MT* vs. initial *ID* for the always Expanding and the always Stationary conditions

For the data illustrated in Figure 2, there was a significant effect of Condition ($F_{1,11} = 12.97, p = .0041$) and *ID* ($F_{3,33} = 305.2, p < .000001$), as well as a significant Condition × *ID* interaction ($F_{3,33} = 12.33, p = .00001$). This interaction reflects the fact, reminiscent of M&B's data, that the *MT* facilitation induced by target expansion increased with the *ID*. A small difference is that, whereas M&B still observed a facilitation effect of well over 100 ms for *ID* = 3.2 bits, in our data Newman-Keuls post-hoc comparisons revealed no significant facilitation for this minimal level of difficulty ($p = .084$), the other pair-wise differences being quite reliable (all three p 's < .002).

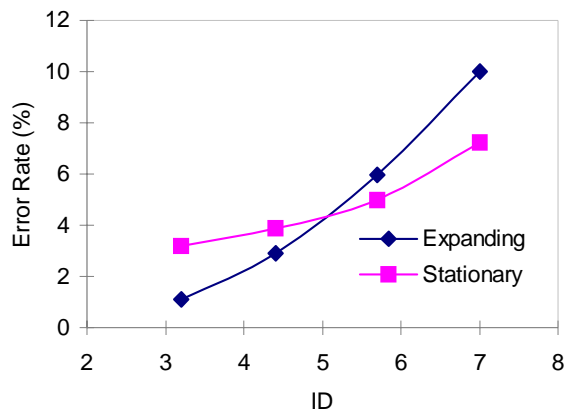


Figure 3. Error rate vs. initial ID for the always Expanding and the always Stationary conditions

Figure 3 shows the mean error rate presented in the same way as *MT* in Figure 2. The only two significant effects here were the main effect of the *ID* ($F_{3,33} = 10.94$, $p = .00004$) and the Condition \times *ID* interaction ($F_{3,33} = 3.58$, $p = .024$). If the monotonic increase of error rate with task difficulty, in parallel with *MT*, is a classic, the interaction pattern illustrated in Figure 3 reveals a potential problem for target expansion. The pattern suggests that while target expansion improved movement accuracy for *ID* = 3.2 bits, it tended to impair accuracy for *ID*s above 5 bits. Post-hoc tests showed that the rightmost pair-wise difference was indeed significant ($p = .019$; for the other three comparisons, $p > .05$), suggesting that target expansion did inflate error rates for the most difficult targets, on average from 7.2% to 10 %.

The data of Figure 3 do not allow us to reach with certainty the same conclusions as M&B, who did not report misses separately. Even though our combined *MT* and error-rate data remain consistent with the view that predictable target expansion facilitates pointing for easy targets, they suggest that with smaller, more difficult targets—obviously those which count in the context of HCI—the speed benefit of expansion was weakened by an accuracy impairment.

Figure 4 shows the distribution of mouse clicks in the two relevant conditions. The *x* axis is normalized, with click position expressed as the distance from target center divided by one half of the target width. Clearly, the distribution of clicks was shaped by the late change of target size. When the final size was doubled, the distribution of hits was flattened, with more points falling outside of the (-1, 1) interval which corresponds to the initial boundaries of the target.

Figure 5 plots the *MT* as a function of the *final* value of the *ID* for the always Expanding and the always Stationary conditions. The figure shows that, at least for higher levels of *ID*, the time needed to reach a target whose size was eventually doubled in the end of cursor

approach was approximately that needed to reach a target that had precisely this final size from the outset. This pattern confirms the participants' ability to perform their movements on the basis of the final size of the target, even when the expansion takes place as late as when 10% of the distance remains to be covered. The reason for such a beneficial effect, however, could be due to either anticipation of or on-line adjustment to the expansion or both.

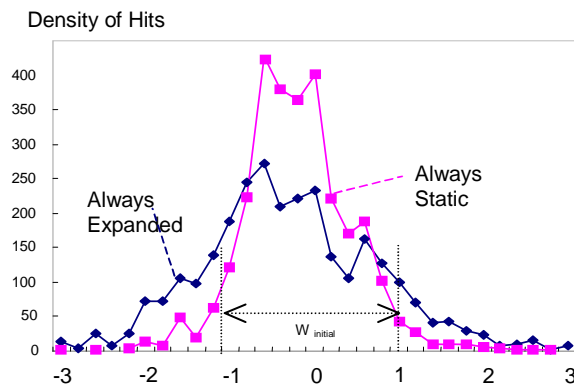


Figure 4. Horizontal distribution of mouse-click normalized by initial target width for the two conditions in which final target size was predictable

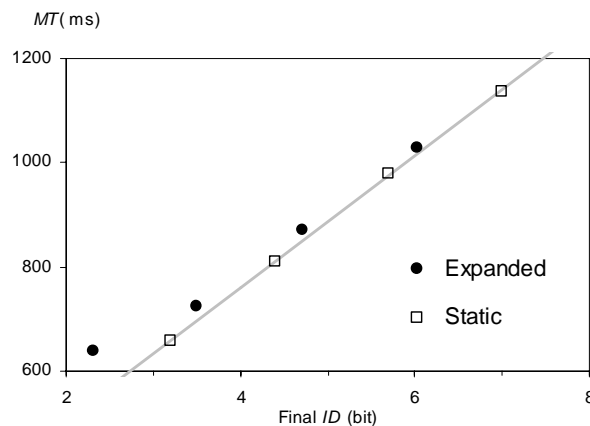


Figure 5. *MT* vs. final *ID* in the two certainty conditions. The regression line has been computed on the basis of the four data points corresponding to the static targets.

Effect of Unpredictable Target Expansion

We now turn to the analysis of performance in the Random condition, in which targets would unpredictably expand, shrink, or remain unchanged after 90% of distance coverage with the cursor. This section focuses on the pair-wise contrast between unchanged and expanding targets.

As shown in Figure 6, expansion caused a shortening of *MT* relative to the unchanged case (for the main effect of Target Type, $F_{1,11} = 69.04$, $p = .000005$), with *MT* being, as usual, strongly and linearly dependent on the *ID* ($F_{3,33}$

= 209.3, $p < .000001$). The ANOVA revealed a significant Target Type \times ID interaction ($F_{3,33} = 8.61$, $p = .0002$), reflecting a consistent reduction of Fitts' law slope when the target unexpectedly expanded (90ms/bit) relative to the unchanged case (118ms/bit)—or, equivalently, a linear increase of the expansion benefit with the scaling up of the ID .

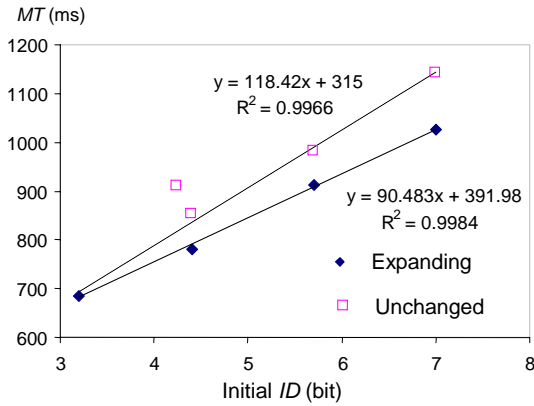


Figure 6. *MT vs. ID for expanding vs. unchanged target size in the Random condition.*

Note that this shortening of MT with the expanded targets firmly establishes that participant could respond online, and in an adaptive way, to target changes that could not be anticipated.

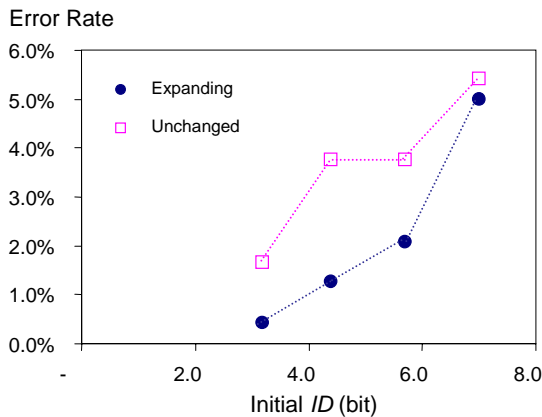


Figure 7. *Error rate vs. ID for expanding vs. unchanged target size for the Random condition*

Turning to movement accuracy, we see in Figure 7 that in our Random condition no accuracy cost was associated with the speed benefit caused by target expansion. Indeed, the mean incidence of errors was arithmetically lower for expanded than unchanged targets, though not reliably so—the main effect of Target Type on error rate fell short of significance ($F_{1,11} = 3.35$, $p = .094$). The only significant effect detected by this ANOVA was that of the ID , whose increase again tended to inflate error rates ($F_{3,33} = 6.91$, $p = .001$).

As visible in Figure 8, the overall effect of target expansion was quite similar in the case of an unpredictable change as in the case of a fully predictable change (Figure 4). The figure also shows that shrinking targets induced a narrowing of the distribution, confirming that, even when more—rather than less—feedback-based correction were demanded, our participants were capable of a substantial degree of on-line adjustment in the face of a rather late modification of target width.

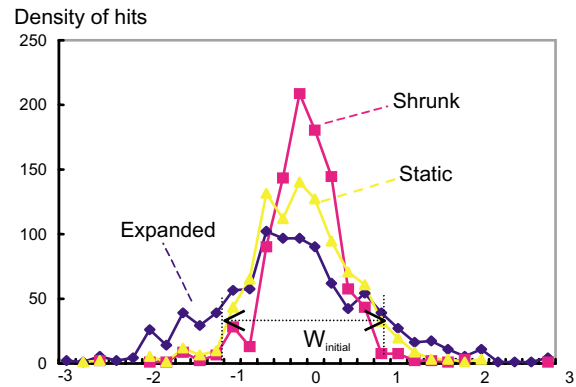


Figure 8. *Distribution of mouse clicks normalized by initial target width in the Random condition*




Figure 9. *MT vs. final ID in the Random condition. The regression line has been computed on the basis of the four data points corresponding to unchanged targets*

Figure 9 plots the MT as a function of the *final* value of the ID for the expanding and unchanged targets of the Random condition. For the more difficult targets, the time needed to reach a target whose size was unexpectedly doubled toward the end of cursor approach was approximately that needed to reach a target that had this final size from the outset. The match, however, becomes mediocre with lower level of final ID . Presumably this is because the faster the movement, the less time available for on-line, feedback-based corrections. At any rate, this pattern confirms quite clearly the participants' ability to

adjust their movements online, even with a very late and inadvertent expansion.

In sum, our data from the Random condition strongly suggest that even when target expansion is made unpredictable by carefully shuffling, within the same block of trials, targets that expand, shrink, and stay unchanged, a rather late expansion of the target does help pointing—specifically, in the case of our participants, by reducing *MT* without increasing the frequency of target misses.

Reaction Time

Reaction times exhibited little variation (around 330 ms) in this study. No significant difference was found between the three conditions (Figure 10). Interestingly, however, the *ID* had a significant impact on *RT* ($F_{3,33} = 13.6$, $p < .0001$), with a pattern reminiscent of that observed for *MT* but with a far lesser magnitude. This finding, which replicates classic reports [6], is a reminder that movement difficulty not only affects the time needed to execute a targeted movement, but also the time needed to prepare it.

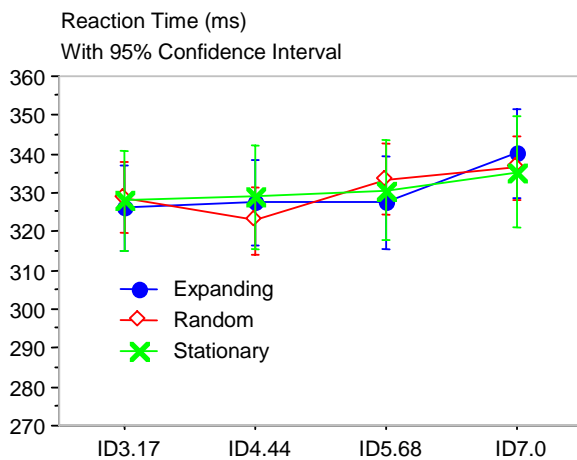


Figure 10. Reaction time for each of the three conditions

Space-Time Correspondence

Figure 11, which shows normalized average kinematic traces for the two certainty conditions, helps understand why participants could still react to a change in target size when 90% of target distance had already been covered. For the Stationary condition, it took only about 55% of *MT* to cover 90% of target distance, thus leaving 45% of the time for the final 10%. In the always Expanding condition, the time spent on the final 10% dropped to 40% of *MT*: target expansion allowed the participants to land faster. In contrast, in the Random condition (Figure 12), participants spent on the last 10% of the distance a larger than normal proportion of *MT* (about 53%), presumably because they had to react to late changes. Anyhow, given the rather large proportion of *MT* (in the range 40%-55%)

allocated to the last 10% of the movement, it is not surprising that late target expansion could facilitate performance—even in the case of an unexpected change.

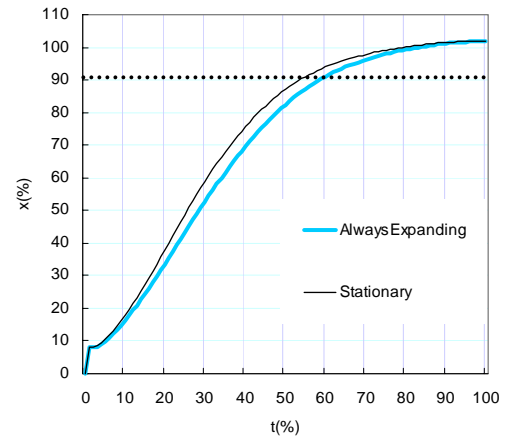


Figure 11. Time expenditure for the always Expanding and the always Stationary conditions (100% of $t = MT$)

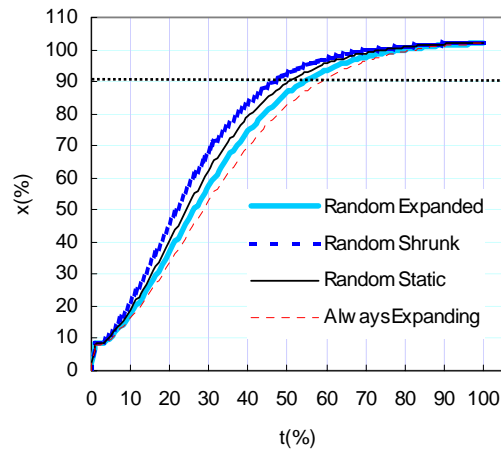


Figure 12. Time expenditure for the Random condition, with the always Expanding condition as a reference

DISCUSSION AND CONCLUSION

Human Response to Target Expansion

Overall, the present experiment offered firm evidence that people could take advantage of late target expansion based on visual information alone without sure knowledge that expansion would occur. However, we obtained less clear-cut evidence of the M&B effect, namely that expansion facilitates pointing in the case of fully predictable expansion, due to the increase of error rate. Although it seems somewhat surprising that expansion leads to a more doubtful advantage when advance information is available, such an outcome is plausible. When the participants were certain about final target expansion, they have anticipated the boundaries of the expanded targets from the outset and ‘aggressively’ taken advantage of it to reduce completion time. Such a speed

advantage could be partially due to visual imagery (expectation of the boundary expansion) and partially due to actual visual information, hence leading to an increased error rate when the targets were small. In the Random condition, in contrast, our participants must have waited until the final change occurred (or failed to occur) to change the characteristic of their current movement, hence the absence of error inflation.

Implications for User Interface Design

Understanding the effects of target expansion on user's behavior and performance is a necessary step in developing interface techniques that take advantage of these effects, but we doubt it is sufficient. The expansion effects could be easily utilized in interfaces if the computer system knew which object is the user's current target. However this is an unreasonable general assumption—if the system knew a priori which object the user is intending to reach, then there would be no need for user's manual interaction at all. So the challenge is to be able to expand the right object without such an assumption.

One encouraging result from the current study is that people could still take advantage of the target expansion even if it happened unexpectedly, as demonstrated in the Random condition of our experiment. This means that prediction algorithms could be used to estimate the probability of an object being the target and to expand it only when such a probability is close to 1.

A particularly relevant case for target expansion is when targets are close to each other, or even packed together. Indeed, if targets are far apart, there is no need to predict which one to enlarge: they can all be enlarged, if only in motor space. We focus the rest of this section on linear arrangements of targets which occur, for example, with the items of a menu, the menus in a menu bar, the tools in a palette, the entries on the MacOS X Dock or Windows task bar, or the items of a list.

A Failed Attempt: The MacOS X Dock

The MacOS X Dock uses a fish-eye view to enlarge the entry under the cursor as an attempt, apparently, to facilitate pointing. Unfortunately, as noted by McGuffin & Balakrishnan (2002), and demonstrated by [7] a fish-eye view provides no pointing advantage. Even though the fish-eye view does enlarge targets in display space, it makes no difference in motor space. As illustrated in Figure 13, as the cursor moves from the left side to the right side of the target, the target shifts from right to left. As a result, the left and right limits for pointing that target are exactly the same as when magnification is disabled. This can be proved formally: if n entries are displayed over p pixels, and if the fish-eye focus only depends on cursor position, there must be a one-to-one mapping between a pixel on the screen and the entry it selects. Since all entries are equal, each entry can occupy only p/n

pixels regardless of the distortion. In order to gain an advantage in motor space as well as display space, magnification of the target cannot be based only on current cursor position.

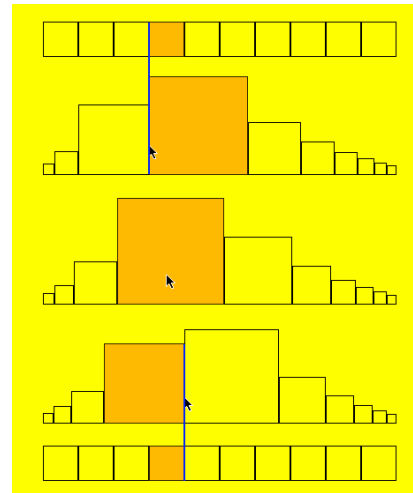


Figure 13. A fish-eye view focused at the cursor position does not change the size of targets in motor space

Note that the MacOS X Dock uses a distortion that actually enlarges the overall size of the Dock when it is magnified. Therefore the targets use a space of $(p+e)$ pixels when distorted rather than p . This results in an advantage of only e/n pixels per target in motor space. For the advantage to be equivalent to doubling the target size, the size of the whole Dock would have to be doubled, making the fish-eye distortion useless.

The Dock Revisited

What is needed is a better predictor of the movement's target than cursor position, one that will provide the system with early information about movement termination. We propose to use the current direction of motion. While the shortest path between the current cursor position and the target is a straight line, movements with a mouse may be more or less curved. Nevertheless, a simple linear regression should lead to an easily controllable interaction, provided it rests on a fast closed loop.

Our design works as follows (see Figure 14). When the cursor is above the Dock and moving at sufficient speed, we compute the intersection of the direction of motion and the baseline of the Dock. If there is such an intersection, it becomes the new focus of the fish-eye view. In order to reduce jerky changes in the display, the coordinates of the focus are smoothed out over time. Magnification of the fish-eye view is designed to take effect only when the cursor is close to the Dock, taking advantage of the experimental result that target expansion is effective even at 90% of the pointing distance. Finally, the position of the focus does not change once the cursor

has entered the Dock for at least one second. This is a critical point since otherwise we would lose the motor space advantage as in the original Dock. After that delay and while the cursor is within the Dock, the focus tracks the cursor position as in the original Dock. This corresponds to a situation where the user is browsing the icons in the Dock: visual magnification is important to recognize the icons, but not helpful to pointing performance.




Figure 14. Dock revisited: the red line shows the direction of motion and the red dot the focus for the fish-eye view

With the original design, the target begins to be magnified when the remaining distance between the cursor and the Dock baseline is about 10% of screen height. When the cursor still has a horizontal component of motion when reaching the level of the Dock, this component of motion induces fish-eye motion, and this involves a cost and no gain: the cost is that the lateral moves of the cursor are amplified, thus perturbing the control [7], and this, as explained above, is in the absence of any gain in terms of pointing facility. With the new design, we eliminate the unwanted fish-eye motion and we gain the advantage of real expansion in motor space. Of course, the effectiveness of this new design needs to be proven with a formal user study.

In conclusion, the current study firmly demonstrates human's ability to respond to and take advantage of target expansion on-line, based on visual feedback alone. Such empirical knowledge provides insights to UI design, as illustrated by our redesign of the Dock. This is but a first step towards improving pointing in user interfaces based on the results of our study.

REFERENCES

1. Accot, J. and S. Zhai. Beyond Fitts' Law: Models for Trajectory-Based HCI Tasks. *Proc. CHI'97*. p. 295-302.
2. Accot, J. and S. Zhai. More than dotting the i's - foundations for crossing-based interfaces. *Proc. CHI 2002*. p. 73 - 80.
3. Baudisch, P., *Drag and Pop*, 2002: <http://www.darmstadt.gmd.de/~baudisch/projects/dragandpop/index.html>.
4. Card, S.K., W.K. English, and B.J. Burr, Evaluation of mouse, rate controlled isometric joystick, step keys and

- text keys for text selection on a CRT. *Ergonomics*, 1978. **21**: p. 601-613.
5. Fitts, P.M., The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 1954. **47**: p. 381-391.
6. Fitts, P.M. and J.R. Peterson, Information capacity of discrete motor responses. *Journal of Experimental Psychology*, 1964. **67**: p. 103-112.
7. Gutwin, C. Improving focus targeting in interactive fisheye views. *Proc. CHI 2002*. p. 267-274.
8. ISO, *ISO 9241-9 Ergonomic requirements for office work with visual display terminals (VDTs) - Part 9: Requirements for non-keyboard input devices (FDIS - Final Draft International Standard)*, 2000, International Organization for Standardization.
9. Kabbash, P. and W. Buxton. The "Prince" technique: Fitts' law and selection using area cursor. *Proc. CHI'95*.
10. MacKenzie, I.S., Fitts' law as a research and design tool in human computer interaction. *Human Computer Interaction*, 1992. **7**: p. 91-139.
11. MacKenzie, I.S., A. Sellen, and W. Buxton. A comparison of input devices in elemental pointing and dragging tasks. *Proc. CHI'91*, p. 161-166.
12. MacKenzie, I.S. and S.X. Zhang. The design and evaluation of a high-performance soft keyboard. *Proc. CHI'99*, p. 25-31.
13. McGuffin, M. and R. Balakrishnan. Acquisition of Expanding Targets. *Proc. CHI 2002*. p. 57-64.
14. Meyer, E.D., *et al.*, Optimality in Human Motor Performance: Ideal Control of Rapid Aimed Movements. *Psychological Review*, 1988. **95**(3): p. 340-370.
15. Schmidt, R.A., *et al.*, Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review*, 1979. **86**(5): p. 415-451.
16. Woodworth, R.S., The accuracy of voluntary movement. *The Psychological Review, Monograph Supplements*, 1899. **3**(2 (Whole No. 13)): p. 1-114.
17. Zhai, S., *On the Validity of Throughput as a Characteristic of Computer Input*, 2002, IBM Research Report RJ 10253, Almaden Research Center: San Jose, California.
18. Zhai, S., B.A. Smith, and M. Hunter, Performance Optimization of Virtual Keyboards. *Human-Computer Interaction*, 2002. **17**(2,3): p. 89-129.