Faster Cascading Menu Selections with Enlarged Activation Areas

Andy Cockburn Andrew Gin

Human-Computer Interaction Lab
Department of Computing Science and Software Engineering
University of Canterbury
{andy, atg19}@cosc.canterbury.ac.nz

ABSTRACT

Cascading menus are used in almost all graphical user interfaces. Most current cascade widgets implement an explicit delay between the cursor entering/leaving a parent cascade menu item and posting/unposting the associated menu. The delay allows users to make small steering errors while dragging across items, and it allows optimal diagonal paths from parent to cascade items. However, the delay slows the pace of interaction for users who wait for the delay to expire, and it demands jerky discrete movements for experts who wish to pre-empt the delay by clicking. This paper describes Enlarged activation area MenUs (EMUs), which have two features: first, they increase the area of the parent menu associated with each cascade; second, they eliminate the posting and unposting delay. An evaluation shows that EMUs allow cascade items to be selected up to 29% faster than traditional menus, without harming top-level item selection times. They also have a positive smoothing effect on menu selections, allowing continuous sweeping selections in contrast to discrete movements that are punctuated with clicks.

CR Categories: H5.2 [User Interfaces]: Interaction styles.

Keywords: Hierarchical menus, cascades, target acquisition, target adaptation, bubble cursors.

1 INTRODUCTION

Selecting items from cascading menus is necessary when using most graphical user interfaces, but it can be error prone and frustrating. When implemented poorly, cascading menus demand a high degree of steering accuracy, and they remain slow to use even when implemented well because of an undesirable design compromise that slows interaction through an explicitly implemented delay.

Figure 1 illustrates part of the problem. If the user wants to select the cascade item 'Header/Footer', the shortest cursor-movement path to the item is the diagonal shown. Unfortunately some menu implementations, such as Java Swing's JMenus, immediately unpost the cascade if the cursor leaves the narrow path across the parent item, which requires the user to back-track the cursor to the parent menu item to re-post the cascade. The common solution, implemented in most widget sets, imposes a short temporal delay between entering/leaving the cascade parent item and posting/unposting its associated menu. The user can pre-empt the delay, which is approximately one-third of a second, by clicking on the parent item. There are also several commercial variations on the delay scheme, including Microsoft Office menus, which postpone the initiation of the delay period until the mouse-velocity falls below a threshold value.

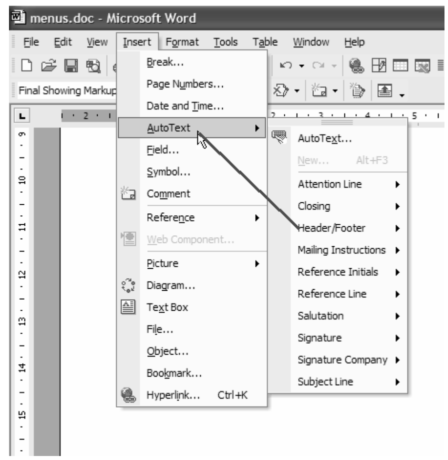

Figure 1. The shortest path in cascade selection is a diagonal, but many menu implementations forbid this path or allow it through a delay that slows interaction.

The delay 'solution', however, is a compromise for three reasons: first, the delay slows the natural pace of interaction, particularly for users who prefer to select items by dragging; second, the delay will be too long for some and too short for others, and individual preferences will change with fatigue and context, limiting the effectiveness of preference settings; third, expert menu users who wish to pre-empt the delay must separate their selection into discrete movements that include a click-to-post before sweeping to the sub-menu.

This paper presents a simple scheme, called Enlarged activation-area MenUs (EMUs), which modifies the behaviour of cascading menus to overcome the disadvantages of temporal delay. Inspired by Bubble-Cursors [10], EMUs increase the activation area for each cascading item so that it includes either the full *y*-dimension of the cascade (Figure 2b) or the maximum *y*-dimension distance to the next cascading menu item (Figure 2c). EMUs also have no posting delay, so the cascade menu appears the instant the cursor enters the activation area. We contend that EMUs are robust and predictable because, unlike other techniques, they are not dependent on subtle temporal parameters.

We evaluate the effectiveness of EMU cascading menus, showing that they reduce selection times by up to 29% with no difference in error rates, and that they have a positive 'smoothing' effect on the way in which users interact with menus.

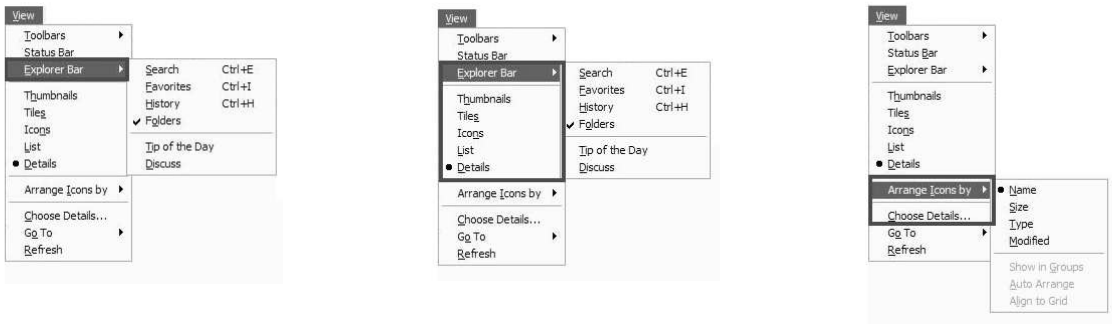

(a) Traditional menus. (b) Full enlarged area. (c) Truncated enlarged area.

Figure 2. Activation areas for cascade menus, as shown by the red box.

2 RELATED WORK

There has been extensive prior research into improving menus, both in developing new interfaces for item selection and in deriving theoretical frameworks for analysing and predicting user behaviour.

There are three main theoretical models that can predict aspects of interaction with menus. First, Fitts' Law [9] predicts the time required to move the cursor to a particular item. Second, Steering Law [1] models the time taken to steer the cursor through a constrained path, such as that used to select cascade items. Third, the Hick-Hyman model of choice-reaction time [11, 12] models the time taken to select one item from a set, and how performance improves with practice. A detailed discussion of these theoretical models is beyond the scope of this paper.

Pie menus [5] were an early attempt to improve menu selection performance by arranging menu items in a circle around the cursor, with each item contributing a 'slice of the pie'. Fitts' Law predicts that pie menus will outperform traditional linear menus because a directional movement of only one pixel is sufficient to reach each item, and because longer movements result in larger targets. Advancing the pie menu concept, Marking Menus [14, 17] allow pie menu items to be selected with a simple directional gesture—the actual pie-menu need not be posted unless the user hesitates in their gesture. Evaluations of pie and marking menus show positive results [14, 17], but their deployment remains limited to niche markets such as advanced graphics/design products.

Split Menus [16] also reduce the Fitts' Law targeting requirements by ensuring that the most frequently accessed items are displayed at the top of traditional menus, above a split in the menu. Findlater and McGrenere [8] further explored these ideas by having split menus adapt to the user's past actions, dynamically rearranging items so that recently selected items are at the top. Their evaluations suggest that adaptive menus harm user efficiency, probably because they inhibit the users' ability to exploit their spatial memory. Kobayashi and Igarashi [13] proposed another distance-reducing menu adaptation, which used the direction of cursor movement to aid posting and unposting cascade menus near to the cursor.

These menu adaptations all aim to improve menu selection times by reducing the distance to the target. Fitts' Law also predicts that acquisition can be improved by increasing the target size. Walker and Smelcer [18] described a 'Fittsizing' approach which statically increased the size of items further from the initial position of the cursor. Fisheye Menus [3] dynamically configure the visual size of menu items, which grow larger as the cursor approaches. Although this technique allows many items to be displayed on the screen at once (eliminating the need for scrolling), the items do not grow in motor-space—although the items are enlarged when the cursor is over them, their actual physical size for targeting remains unchanged. The evaluation of Fisheye Menus showed them to be slower than traditional cascading menus [3]. Rather than enlarging the targets, bubble cursors [10] dynamically alter the size of the cursor's activation area so that it always encloses at least one target. Evaluations showed that bubble-cursors are efficient for abstract targeting tasks. We are unaware of prior practical applications of the technique.

Experiments with multi-modal feedback for aiding menu item acquisition have shown mixed results. Brewster and Crease [4] showed that audio feedback can reduce errors in menu selection, and that it allows errors to be identified more quickly. Dennerlein, Martin and Hasser [7] showed that force-feedback can help in steering tasks, but the results of Oakley, Brewster and Gray [15] suggest that when targets are not discrete (as is the case in menus), force feedback can harm selection times. Cockburn and Brewster [6] evaluated all combinations of audio, tactile and 'sticky' widgets in both discrete and non-discrete conditions, using menus for the non-discrete condition. Results showed that all tested modes of feedback aided acquisition with discrete targets, but that the benefits are unclear in menus, with stickiness harming menu selection times significantly.

Finally, Ahlström [2] modeled and improved cascading menu selection times through the use of 'force-fields', a variant of sticky widgets, that attracts the cursor towards the cascading menu. The evaluation did not investigate whether the technique caused an adverse effect on selecting non-cascading items.

3 ENLARGED ACTIVATION AREA MENUS: DESCRIPTION & PILOT STUDY

To recap the problems with traditional menu cascades: shortcut paths or steering errors will cause cascade menus to be unposted unless there is a delay, but delays slow the pace of interaction; and expert users who pre-empt the delay by clicking must use discrete target acquisitions, punctuated by stationary clicks.

EMUs aim to overcome these problems by eliminating the delay and by enlarging the activation area of the parent menu associated with each cascade. The enlarged area includes all of the menu items below the cascade-parent up to the end of the cascade menu (Figure 2b) or to the next cascade parent (Figure 2c), whichever comes first. The enlarged activation area allows users to slip off the parent item or to take a diagonal path to the cascade target without the use of delays.

There are two potential limitations of enlarged activation areas:

1. The effectiveness of the technique partially depends on the density of cascading items in the parent menu. The relative positioning of the cascading items also impacts on their potential efficiency. Each activation area extends at most to the next cascading parent (Figure 2c). Therefore, if every item in the parent menu is a cascade, then none of the activation areas will be enlarged. However, even in this worst-case scenario, it is possible that the removal of delay will, on average, improve performance.
2. It is possible that users will be confused or distracted by the appearance of a cascade menu when they are targeting a non-cascade item that lies within the enlarged activation area. This seems unlikely because users are already accustomed to cascade menus 'lingering' while they point at other items due to the current delay strategy.

To partially quantify the concerns about menu-cascade density and positioning, we inspected top-level cascading menus in the default-install state of three common desktop applications. Only menus with at least one cascade were included in the analysis. Table 1 summarises the results for Microsoft Word, Microsoft Internet Explorer and Adobe Reader. The average density of cascade entries (the number of cascade items divided by the total number of items) across these three interfaces is 0.26 (sd 0.23)—meaning that in menus that contain cascades, on average around one in four top-level items is a cascade entry. The other metric influencing the potential efficiency of the technique is the proportion of the enlarged activation area that is available for use—if, for example, there is only one cascade item in a menu, then 100% of the area is available, but if the item below the cascade item is also a cascade item, then 0% of the enlarged area (the area in addition to the entry itself) is available. Across these three interfaces, an average of 67% of the enlarged area is available (sd 32%). These data suggest that enlarged activation areas are viable within current commercial software systems.

3.1 First Prototype and Pilot Study

We initially experimented with enlarged activation areas by modifying the behaviour of Java Swing JMenu objects. The system supported two menu modes—traditional JMenus and enlarged activation areas, both of which also used the JMenu default delay of 200ms.

In the pilot study, 14 volunteer participants (all right-handed male post-graduate students) used both traditional JMenus and EMUs. The order in which the participants used the two menu types was counter-balanced, with half of the participants completing all JMenu tasks first, and half using EMUs first. All target items were contained within a cascade menu that was the fourth entry in the top-level menu. The target items and their parents were highlighted green. The targets were either the first, fourth or seventh item in a seven-item cascade. Participants were allowed around one minute to practice with each menu type before completing three selections of each of the three items (1^{st}, 4^{th} and 7^{th}). Having completed all 18 selections (2 menu-types, 3 distances, 3 repetitions), the participants completed to a short preference questionnaire.

Software automatically controlled exposure to the experimental conditions, and it logged all task times and errors. The Java experimental program is available at www.cosc.canterbury.ac.nz/~andy/menus/menutester.jar.

Data was analysed using a 2×3 repeated measures analysis of variance (ANOVA) for factors *menu-type* (JMenus, EMUs) and *target-location* (1^{st}, 4^{th} and 7^{th} item in the cascade), with subject as a random factor. We hypothesized that EMUs would allow faster menu selection, and that they would become comparatively more efficient as target items are lower in the menu because they allow more direct diagonal paths.

Table 1. The density of cascading menu items in top-level menus that contain at least one cascade item for the install-state of Microsoft Word, Internet Explorer, and Adobe Reader. Also, the percentage of the enlarged activation area available.

	Density mean (sd)	% enlarged area available mean (sd)
Word	0.16 (0.11)	76 (23)
IE	0.29 (0.15)	63 (32)
Reader	0.34 (0.34)	60 (43)
Overall	**0.26 (0.23)**	**67 (32)**

3.2 Results

Eleven of the 252 trials (4.4%) resulted in incorrect item selections: five with JMenus, and six with EMUs. Data from these incorrect trials are discarded.

There is a significant main effect for menu-type, with selections from EMUs being 14% faster than JMenus: JMenu mean 1.47s (sd 0.4), EMU mean 1.26s (sd 0.2), $F_{1,13}=7.3$, $p<.05$. As expected, there is a significant main effect for target-location, with the distance to target influencing performance: $F_{2,26}=20.4$, $p<.01$. There is also a significant menu-type×target-location interaction ($F_{2,26}=7.9$, $p<.01$), with, as hypothesized, EMUs providing greater performance advantages when the target item is lower in the menu.

Subjective responses also favour EMUs. Participants rated their 'liking' and 'efficiency' with the two menu types, using five-point Likert-scales (1 being 'not liked' or 'inefficent' and 5 being 'liked' or efficient'). Median responses were both 5 for EMUs and 3.5 (liking) and 3 (efficiency) for JMenus, showing significant differences (Wilcoxon matched-pairs, p<.01).

4 MAIN EVALUATION

While encouraging, there are three clear limitations in the pilot study. First, JMenus do not implement an unposting delay, failing to reflect current best-practice in menu behaviour. Second, the experiment only inspected selection from cascade items, but improvements in cascade-selection performance are only profitable if they cause no adverse effects when selecting non-cascading items. Third, enlarged activation areas allow the posting delay to be eliminated, but this delay was retained in the pilot study.

We therefore conducted a second study of enlarged activation areas. To allow complete control over menu behaviour, we implemented cascading menus from elementary graphics operations using Tcl/Tk. The experimental system is accessible from www.cosc.canterbury.ac.nz/~andy/menus/menus.tcl.

Tasks either involved selecting the 1^{st}, 4^{th} or 7^{th} item from a cascade menu associated with the 3^{rd} item in the parent menu, or they involved selecting a top-level item that was the 4^{th} item in the parent menu. Targets and their parents were identified by red highlighting. The inclusion of top-level selection tasks allows us to inspect whether any of the conditions harm selection of non-cascading items. Data from the cascade selections and from top-level menu selections are analysed separately, as described below.

4.1 Experimental Design

The experiment is primarily designed to expose differences in cascade selection times across the factors that influence the effectiveness of enlarged activation areas. Data from the cascade

(a) Medium density (b) 'Dense' density

Figure 3. The experimental interface cuing the selection of the 4th sub-item of the cascade in the 'medium' (left) and 'dense' (right) conditions. The title-bar shows the current expansion and delay configuration.

item selections are analysed in a 2×2×2×3 repeated measures ANOVA (with subject as a random factor) for the following factors:

- *activation area*—normal or enlarged;
- *delay*—0 or ⅓s (333ms) posting/unposting delay;
- *density*—medium or dense cascade items in the parent menu;
- *target-location*—either 1st, 4th or 7th cascade item.

This design allows us to separately inspect the efficiency contributions of the two main design features of EMUs: the size of the activation area and the length of the posting/unposting delay. Also, the density factor allows us to inspect how performance is influenced by the proximity of nearby cascade items in the parent menu. In the 'medium' density condition (Figure 3a), three of the eleven parent items were cascades, allowing 34% of the enlarged area to be used to benefit selection. In the 'dense' condition (Figure 3b), ten of the eleven items were cascades, allowing 17% of the enlarged activation area to be used (the 4th item in the parent menu was always a non-cascading item). Finally, the target location factor should reveal differences across conditions for targets at different locations in the cascade menu.

Data from the top-level item selections are analysed using a 2×2×2 repeated measures ANOVA for factors activation area, delay and density. We anticipate that none of these factors will significantly influence top-level selection times—the factors affect the appearance of the cascade menu, but they should not impact on selecting the 4th item in the parent menu—unless the participants are confused or distracted by the cascade-posting in the enlarged activation area.

4.2 Participants and Procedure

Fourteen right-handed male Computer Science graduate students took part in the experiment, three of whom had participated in the pilot study two months earlier. Each participant completed all tasks with either normal or enlarged activation areas before proceeding to the other condition. This order was counter-balanced between participants.

Participants completed six blocks of tasks with each activation area setting, including two initial unanalyzed training blocks. The four logged blocks covered the conditions: medium density and zero delay; medium density and ⅓s delay; dense and zero delay; and dense and ⅓s delay. The window title-bar displayed the particular expansion and delay setting for each block (Figure 3). At the start of the experiment the participants were instructed to take note of the cue prior to commencing the tasks in each block. The order of the logged blocks was randomised for each participant. The two training blocks had only one cascading item in the parent menu. Participants were encouraged to rest their fingers and wrists between blocks, and software enforced a rest of at least five-seconds.

Each block consisted of thirteen tasks: two initial unanalyzed preparation tasks, then eleven analysed tasks comprised of three repetitions of each target-location (1st, 4th or 7th cascade item) and two selections of top-level item 4. The order of the non-preparation tasks within each block was randomized to reduce anticipation.

Incorrect selections were logged, but the software continued to cue for the same task until successfully completed. In the results analysis we report the number of errors, but task times are reported from the last time that the top-level menu was posted prior to successful acquisition. We also analysed the data from the initial posting time, which shows exactly the same main-effects and interactions being significant.

4.3 Apparatus

The Tcl/Tk interface controlled the participants' exposure to all conditions, and it logged all user actions including all mouse movements within the menus.

The experiment ran on Intel Pentium 4 2.8GHz computers running Fedora Core Linux. The computer was equipped with 1GB RAM, and connected to a 19-inch CRT display at a 1600x1200 resolution and a 75Hz refresh-rate. The interface window was 500x700 pixels. Each menu item was 22 pixels high. The top-level menu was 100 pixels wide, and the cascade was 120 pixels wide (Figure 3). Input was received through a Logitech three-button wireless-mechanical mouse set to the default control-display gain.

5 RESULTS

We present the results in the following order. First, we describe the main results concerning cascade menu item selection. Second, we directly compare cascade selection with traditional menus against the intended design of EMUs. Third, we analyse the data top-level menu item selection. Finally, we characterise the mouse-movement patterns.

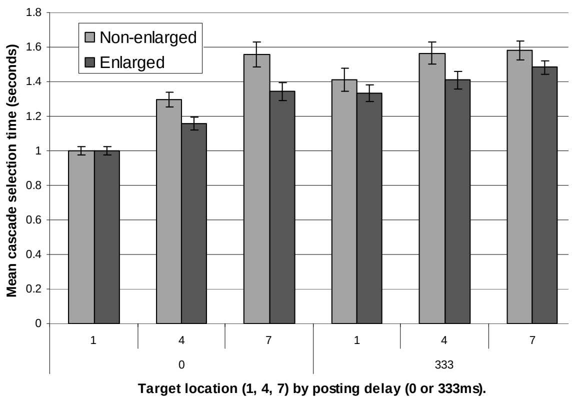

Figure 4. Mean time for cascading item selection for normal and enlarged activation areas by target location and posting delay.

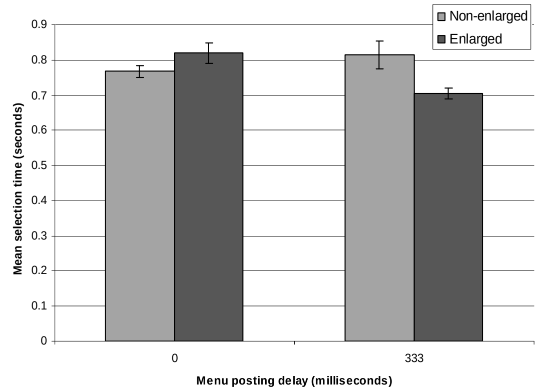

Figure 5. Mean selection times for selecting a non-cascading item from the top-level menu.

5.1 Cascade Selections

The mean time for the 1008 cascade item selections across all conditions was 1.35s (sd 0.33). There were a total of 17 selection errors: three with enlarged areas and zero delay; three with enlarged areas and ⅓s delay; three with normal areas and zero delay; and eight with normal areas and ⅓s delay.

There is a significant difference between the mean times to select items with normal activation areas (mean 1.4s, sd 0.36) and enlarged activation areas (mean 1.29s, sd 0.27): $F_{1,13}$=30.5, p<.01. There is also a significant main-effect for factor *delay* ($F_{1,13}$=23.8, p<.01), with a zero delay mean of 1.23s (sd 0.31) and a ⅓s-delay mean of 1.46s (sd 0.3). The main effect for *target-location* shows strongly significant differences ($F_{2,26}$=72.9, p<.01): 1st item 1.19s (sd 0.3), 4th item 1.36s (sd 0.3), and 7th item 1.49s (sd 0.3). Finally, the *density* of the cascade items in the parent menu showed no significant main effect (F<1, ns), with very similar selection times for the two levels.

There is a significant interaction between factors *delay* and *target-location*: $F_{2,26}$=18.4, p<.01. The interaction, apparent in Figure 4, is caused by the participants' relatively rapid performance deterioration in the zero delay condition as they acquire targets that are located lower in the menu. This interaction is probably caused by participants cautiously steering across parent items to avoid an immediate unposting of the cascade when there is no delay.

No other interactions were significant, but it is worthwhile noting that, unlike the pilot study, the anticipated interaction between *activation-area* and *target-location* was not significant ($F_{1,13}$=2.5, p=0.1). This is probably best explained by the fact that the *activation-area* factor in this study consists of not only the 'intended' EMU design of enlarged areas with zero delay, but also enlarged areas with a ⅓s delay.

5.2 Traditional versus Enlarged Activation Areas

The analysis above separates the two interface properties that comprise the intended behaviour of EMUs: zero delay and enlarged areas.

To directly compare traditional cascade selections with EMUs, we planned to compare the data from the ⅓s delay+no-enlargement condition with data from the zero-delay+enlargement condition, using a 2×3 ANOVA for factors *interface-type* (traditional versus EMUs) and *target-locations* (1st, 4th, 7th).

As expected, both main effects show significant differences—interface-type $F_{1,13}$=38.1, p<.01; target-location $F_{2,26}$=32.9, p<.01. The important result, however, is that mean performance with EMUs (1.17s, sd 0.22) is 23% faster than traditional menus (1.52s, sd 0.3): ranging from 29% faster for selecting the 1st item to 15% faster for selecting the 7th item. The interface-type×target-location interaction was not significant: $F_{2,26}$=2.6, p=.09.

To check that the positive results for EMUs were not caused by a few participants using them particularly well, we calculated per-participant means with EMUs and traditional menus. Every participant was faster with EMUs, with the individual's performance benefits ranging from 7% to 39% faster.

5.3 Top-level item selection

Performance improvements in selecting items from cascades are only profitable if there is no harm to selecting non-cascading items from the parent menu. To recap, data from the top-level selections are analysed using a 2×2×2 repeated measures ANOVA for factors *activation area*, *delay* and *density*.

The mean time for top-level item selections was 0.78s (sd 0.14). There were 15 incorrect selections from 224: three with enlarged areas and zero delay; nine with enlarged areas and ⅓s delay; none with normal areas and zero delay; and three with normal areas and ⅓s delay. The nine errors with enlarged areas and ⅓s delay are a concern, and their cause is unclear, but fortunately, the intended design of EMUs does not couple enlarged areas with non-zero delays.

None of the factors showed significant main effects. The mean selection time with enlarged areas was 0.76s (sd 0.1) compared to 0.78s (sd 0.2) without: $F_{1,13}$<1, ns. The only significant interaction is caused by a cross-over effect, shown in Figure 5, between factors *delay* and *activation areas* ($F_{1,13}$=8.8, p<.01). When using normal areas, mean performance deteriorates slightly when the delay increases from zero to ⅓s, but when using enlarged areas mean performance improves between zero and ⅓s delays. This effect suggests that participants found the immediate posting of the enlarged areas slightly distracting to their task of top-level item selection.

5.4 Characterising Movement-to-Target

To help compare and characterise the participants' use of the different menu conditions, we developed a tool that displays the low-level mouse actions used to acquire cascade items. Figure 6, for example, shows the paths taken by participant 8, for factors *activation-area* and *delay*, with enlarged areas in the right, column and zero delay on the top row. Figure 6b shows that

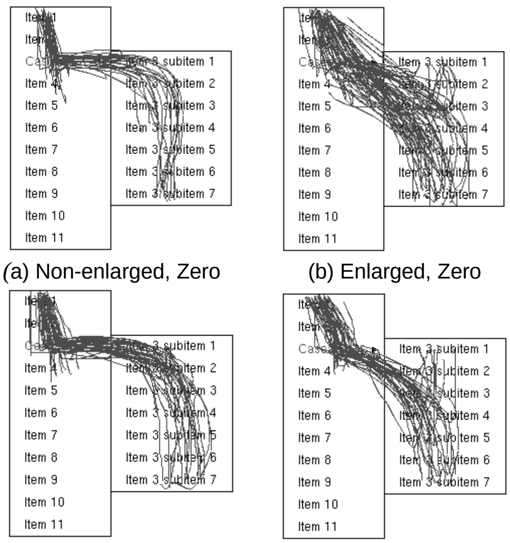

(a) Non-enlarged, Zero (b) Enlarged, Zero

(c) Non-enlarged, 333 (d) Enlarged, 333.

Figure 6. Cascade selection paths for non-enlarged (left) and enlarged (right) activation areas, and for zero (top) and ⅓s delays (bottom).

EMUs (zero delay, enlarged) result in a fluid sweeping action to select items. In contrast, non-enlarged areas result in two discrete movements (Figures 6a,c), with the first leading to the cascade parent item, followed by an abrupt direction change before carefully steering across the cascade to the item. Figure 6d shows that when enlarged areas are combined with a ⅓s delay, the participant paused on the cascade item (to pre-empt the delay with a click) before sweeping towards the target.

6 DISCUSSION

To summarise the results, enlarged activation areas and zero delays both improve the selection times of cascading menu items. When used in combination, as intended for the EMU design, they improve selection times by up to 29% in comparison to traditional cascading menus. These performance improvements are achieved without harming the selection of non-cascading menu items. Finally, EMUs result in smoother, fluid paths to cascading entries.

Subjective responses also favoured EMUs, with the pilot study confirming stronger perceived efficiency and liking for EMUs than traditional JMenus. Although we did not formally gather subjective assessments in the main experiment, several participants stated that they found the combination of no-delay and enlarged areas to allow ‘simpler’, ‘faster’, ‘more natural’ and ‘less frantic’ menu selections. The path visualisations support these subjective assessments.

There are several limitations in this study, as well as research opportunities, that we wish to address in further work.

Understanding density and neighbour effects. In the main experiment the ‘medium’ density condition had two non-cascading items below the target cascade-parent item (Figure 3); the ‘dense’ condition had one. The absence of a significant performance difference between medium and dense conditions suggests that EMUs are robust to variations in density, provided there is at least one non-cascading item below the target cascade. In future work we will scrutinize this possibility using greater statistical power.

A related issue that we will also examine in further work is how performance with EMUs deteriorates when cascade parent items are immediate neighbours. In this case, it is possible that the zero delay will harm acquisition because users must accurately steer across the item to reach the correct cascade menu. The current scheme, in contrast, allows diagonals on immediately neighbouring items, but at the cost of a delay.

Systematic variation of parent-item location. The selection tasks in the experiment always involved selecting an item from a cascade menu associated with the third item in the parent menu. We wish to confirm that EMU performance benefits are robust across all menu locations. Similarly, we would like to confirm that the benefits are retained when selecting from 2nd level cascading menus. We currently see no reason to suspect that this should not be the case.

Variation and combination of cascade menu behaviours. There are a variety of subtly different behaviours for cascading menus demonstrated by current desktop environments. For example Mozilla Firefox/Thunderbird use a reciprocal posting/unposting delay, while Microsoft Office applications postpone the unposting delay until the cursor velocity falls below a threshold value. To improve our understanding of human performance with cascades we intend to experiment with a variety of features, including non-reciprocal delays, velocity-based delay postponement, and enlarged areas. As well as investigating maximal performance with these different techniques, we also wish to inspect their robustness to changes in the user’s context of work, such as when first learning the menu structure and when using mobile input devices such as track-pads.

7 CONCLUSION

We have presented enlarged activation area menus that are designed to aid the selection of items contained in cascading menus. The work was inspired by bubble-cursors, which dynamically adapt the cursor’s activation area so that it always encloses exactly one target. The enlarged area associated with each cascade menu allows our technique to eliminate the explicit delay that is used in current implementations prior to posting and unposting cascading menus. Evaluations show that enlarged activation areas and zero delays improve cascade-item selection by up to 29% in comparison to traditional methods, and that it has a positive ‘smoothing’ effect on acquisition paths.

ACKNOWLEDGEMENTS

Many thanks to the anonymous reviewers who provided extremely helpful and detailed comments.

REFERENCES

[1] Accot, J. and Zhai, S. Beyond Fitts' Law: Models for Trajectory-Based HCI Tasks. in Proceedings of CHI'97 Conference on Human Factors in Computing Systems (Atlanta, Georgia, 1997), 295--302.

[2] Ahlström, D. Modeling and Improving Selection in Cascading Pull-Down Menus Using Fitts' Law, the Steering Law and Force Fields. in Proceedings of CHI'05 Conference on Human Factors in Computing Systems, (Portland, Oregon, 2005), ACM Press, 61-70.

[3] Bederson, B. Fisheye Menus. in Proceedings of UIST'00 Symposium on User Interface Software and Technology, (San Diego, California, 2000), ACM Press, 217-225.

[4] Brewster, S. and Crease, M. Making Menus Musical. in Proceedings of INTERACT'97: the sixth IFIP conference on Human Computer Interaction, (1997), Chapman & Hall, 389-396.

[5] Callahan, J., Hopkins, D., Weiser, M. and Shneiderman, B. An Empirical comparison of Pie Versus Linear Menus. in Proceedings of CHI'88 Conference on Human Factors in Computing Systems, (Washington, DC, 1988), ACM Press, 95-100.

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

[7] Dennerlein, J., Martin, D. and Hasser, C. Force-Feedback Improves Performance for Steering and Combined Steering-Targeting Tasks. in Proceedings of CHI'00 Conference on Human Factors in Computing Systems, (The Hague, The Netherlands, 2000), ACM Press, 423-430.

[8] Findlater, L. and McGrenere, J. A comparison of static, adaptive, and adaptable menus. in Proceedings CHI'04 Conference on Human Factors in Computing Systems, (Vienna, Austria, 2004), ACM Press, 89-96.

[9] Fitts, P. The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement. Journal of Experimental Psychology, 47. 1954. 381-391.

[10] Grossman, T. and Balakrishnan, R. The bubble cursor: enhancing target acquisition by dynamically resizing of the cursor's activation area. in Proceedings of CHI'05 Conference on Human Factors in Computing Systems, (Portland, Oregon, 2005), ACM Press, 281-290.

[11] Hick, W. On the rate of gain of information. Quarterly Journal of Experimental Psychology, 4. 1952. 11-36.

[12] Hyman, R. Stimulus information as a determinant of reaction time. Journal of Experimental Psychology, 45. 1953. 188-196.

[13] Kobayahi, M. and Igarashi, T. Considering the direction of cursor movement for efficient traversal of cascading menus. in Proceedings of UIST'03 Symposium on User Interface Software and Technology, (Vancouver, Canada, 2003), ACM Press, 91-94.

[14] Kurtenbach, G., Sellen, A. and Buxton, W. An Empirical Evaluation of Some Articulatory and Cognitive Aspects of Marking Menus. Human-Computer Interaction, 8 (1). 1993. 1--23.

[15] Oakley, I., Brewster, S. and Gray, P. Solving Multi-Target Haptic Problems in Menu Interaction. in Extended Abstracts of CHI'01 Conference on Human Factors in Computing Systems, (Seattle, WA, 2001), ACM Press, 357-358.

[16] Sears, A. and Shneiderman, B. Split menus: Effectively using selection frequency to organize menus. ACM Trans. Comput.-Hum. Interact., 1 (1). 2004. 27-51.

[17] Venolia, D. and Neilberg, F. T-Cube: A Fast, Self-Disclosing Pen-Based Alphabet. in Proceedings of CHI'94 Conference on Human Factors in Computing Systems, (Boston, MA, 1994), ACM Press, 265-270.

[18] Walker, N. and Smelcer, J. A Comparison of Selection Times from Walking and Pull-Down Menus. in Proceedings of CHI'90 Conference on Human Factors in Computing Systems, (Seattle, Washington, 1990), ACM Press, 221-226.