

Comparing Automatic and Manual Zooming Methods for Acquiring Off-Screen Targets

Joshua Savage[†] & Andy Cockburn[‡]

[†]*LeftClick Ltd. Canterbury Innovation Incubator. PO Box 13761, Christchurch, New Zealand*

[‡]*Human-Computer Interaction Lab, Department of Computer Science, University of Canterbury, Christchurch, New Zealand*

Tel: +64 3 364 2987

Fax: +64 3 364 2569

Email: Josh.Savage@leftclick.co.nz
andy@cosc.canterbury.ac.nz

Abstract

Previous studies indicate that user performance with scrolling can be improved through Speed-Dependent Automatic Zooming (SDAZ), which automatically couples the document's zoom-level with scroll-speed. These studies have compared traditional scrolling techniques (scrollbars and rate-based scrolling) with SDAZ, leaving a potential confound that the efficiency gains are due to zooming rather than the automatic binding of zoom-level with speed. It is therefore possible that decoupling zoom from speed, allowing users separate but concurrent control of each, could further enhance performance. This paper describes an experiment (n=35) that examines user performance, workload, and preference in tasks that involve scroll-based acquisition of off-screen targets using SDAZ and manual zooming. Three different types of document navigation are explored: text documents, 'flat' 2D maps, and a 'globe browser' that allows multi-level zooming of a globe-map of Earth and underlying city views. Results show that automatic zooming not only improves performance, but that it does so with substantially less subjective workload, and that it is strongly preferred. We also confirm limited previous work using Fitts' Law as a model for off-screen target acquisition, and show that it applies even when zooming is employed.

Keywords: Speed-dependent automatic zooming, scrolling, zooming, target acquisition, Fitts' Law.

1 Introduction

Scrolling is the main interface technique for navigating through documents that are too large to be displayed within a single window. The fundamental importance of

scrolling has led contemporary software and hardware vendors to develop a wide range of enhanced scrolling techniques. These include rate-based scrolling (which is activated by dragging the middle mouse button in Windows platforms), semantic scrolling (supported by Microsoft Word through the extension at the bottom of the vertical scrollbar), mouse-wheel scrolling, and isometric input devices such as the IBM TrackPoint. Variable magnification zooming is another commonly supported interface control that is related to scrolling because it changes the proportion of the document visible within each window. The “Dynamic Zoom” feature of Adobe Reader 6 attempts to exploit the scroll-zoom relationship by allowing concurrent scroll-zoom actions.

When scrolling for target acquisition (such as browsing for a particular heading, picture, or other landmark in a document), it is most efficient to scroll as quickly as possible to the target. But rapid scrolling induces motion-blur [Burr 1980]: the information moves across the screen so quickly that our eyes cannot keep up. Igarashi and Hinckley [2000] proposed ‘speed-dependent automatic zooming’ (SDAZ) as a method to overcome motion blur. With SDAZ, the zoom-level is automatically adjusted as the scroll-rate increases, allowing rapid document movement at visually manageable pixel movement rates (see Figure 1). The scroll-speed is controlled through normal rate-based scrolling—the further the user drags the middle button, the faster they scroll—but automatic zooming means that the faster they scroll, the ‘higher’ they fly above the document.

Igarashi and Hinckley’s preliminary evaluation indicated that SDAZ allows comparable performance to other scrolling techniques. Our earlier work showed that in text-document and map navigation domains SDAZ can outperform commercial systems using traditional scrollbars and panning [Cockburn and Savage 2003]. More recently we described theoretical and empirical measures for calibrating the relationship between speed and zoom, and we eliminated the possibility that our previous results were confounded by rate-based controls outperforming traditional scrollbars [Cockburn, Savage and Wallace 2005]. The evaluation showed that SDAZ outperforms traditional scrollbars, rate-based scrolling, and another variant of automatic zooming based on van Wijk and Nuij’s [2004] work on optimal pan-zoom trajectories.

This paper adds three further pieces to a strong argument in favour of commercial deployment of SDAZ. First, it reports the results of an experiment ($n=35$) comparing off-screen target acquisition using SDAZ (rate-based scrolling plus automatic zooming) with that of rate-based scrolling plus manual zooming. The central hypothesis is that automatic zooming allows faster target acquisition with lower cognitive effort than manual zooming. Second, it extends the prior evaluations, which focused on text-document scrolling, to other domains with maps and a globe browser. Third, it validates a small set of research on using Fitts’ Law as a performance model for off-screen target acquisition [Guiard, Beaudouin-Lafon, Bastin, Pasveer and Zhai 2004; Hinckley, Cutrell, Bathiche and Muss 2002]. Although Fitts’ Law is well known for modelling on-screen target acquisition, we confirm its accuracy for off-screen scroll-based target acquisition, even when zooming is employed.

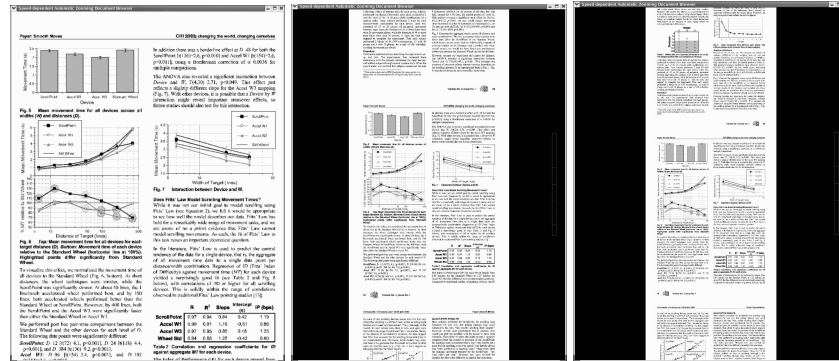


Figure 1. The automatic-zooming text interface with slow, medium, and high speed scrolling.

Further details on background studies are presented in Section 2. Section 3 then describes the three interfaces examined in our evaluation. Section 4 describes the evaluation method, Section 5 presents the results, and Sections 6 and 7 discuss implications and conclude.

2 Background

Although SDAZ was introduced to the research community by Igarashi and Hinckley in 2000, a similar concept was first demonstrated in the computer game ‘Grand Theft Auto’ in 1997. The game gave users a plan view of their car in a city street, which automatically zoomed to show progressively more city blocks on acceleration. The need for zooming in the game is clear: without it, the rate of display change (the speed of pixel movement) can exceed human visual processing limits, inducing ‘motion blur’. Zooming-out decreases the rate of pixel movement, allowing higher speeds in the information space without overloading the visual system.

As described in the introduction, there have been three main evaluations of scrolling interfaces that automatically zoom. First, Igarashi and Hinckley’s preliminary study with seven participants found no definitive performance differences between SDAZ and normal scrolling. Second, in our prior work, we showed that SDAZ allowed users to complete map and text-document browsing tasks more rapidly than traditional scrollbar navigation in standard commercial systems. This result had a potential confound because the experiment compared rate-controlled scrolling (with automatic zooming) against scrollbar scrolling. To eliminate this potential confound, our recent work compared user performance with normal scrollbars, with rate-based scrolling, and with two versions of automatic zooming: one based on manipulation of a virtual scroll-thumb, and the other based on rate-based input. Again, the results favoured the SDAZ behaviour of rate-based input with automatic zooming.

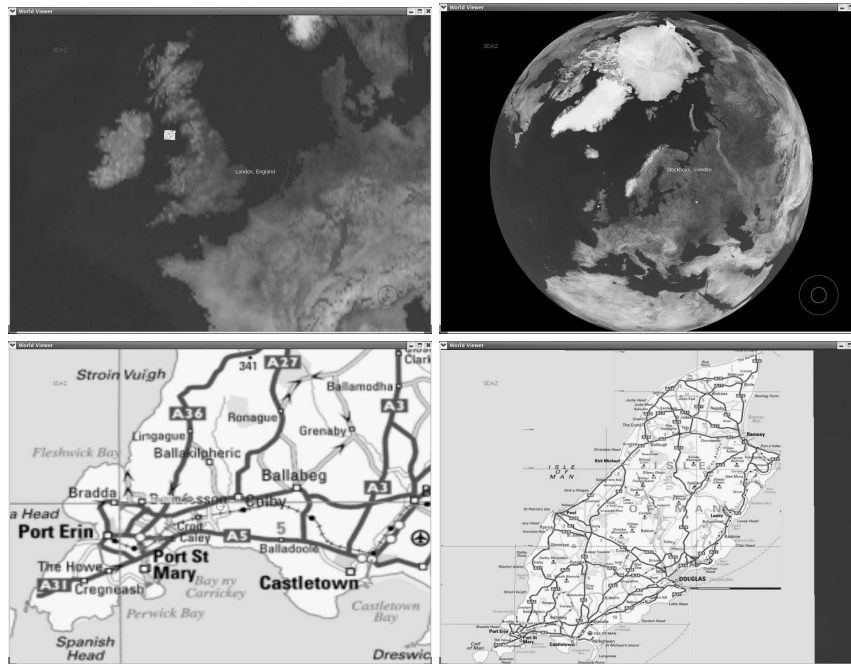


Figure 2. The automatic-zooming globe interface at slow (left) and high (right) speeds in the globe-view (top), and at slow and high speeds in the city view (bottom).

The experiment reported in this paper focuses on the cause of the efficiency improvements of SDAZ. The previous evaluations have shown that automatic zooming allows enhanced performance, but it remains unclear whether automatic zooming is better or worse than manual zooming. It is reasonable to suspect that manual zooming could outperform automatic zooming because it decouples scroll-speed from zoom level, allowing greater independent control of speed and zoom. Furthermore, there is evidence that parallel input of separate controls through bimanual interaction can enhance performance over serial input. Leganchuk, Zhai and Buxton [1998] and Casalta, Guiard and Beaudouin-Lafon [1999] both showed performance benefits for bimanual interaction in rectangle editing tasks. In a domain more closely related to our work, Zhai *et al* [1997] showed that scrolling and pointing tasks are improved by using parallel control separation with a mouse in one hand and a joystick in the other. Finally, Hinckley, Czerwinski and Sinclair [1998] describe and theoretically evaluate two-handed interaction for panning, zooming and rotation, but they did not empirically validate their findings.

2.1 Fitts' law

Fitts' Law [1954] accurately models the time taken to acquire on-screen targets in graphical user interfaces across a very wide range of input devices. The "Shannon

formulation” of Fitts’ Law [MacKenzie 1992] predicts that cursor movement time MT increases linearly with the Index of Difficulty (ID), which is the logarithm of the distance moved (amplitude, A) over the target width (W): $MT=a+b \times ID$, where $ID=\log_2(A/W+1)$; also $IP=1/b$. The constant b , determined through linear regression, provides a useful estimate of hand-eye coordination using the targeting method, and its reciprocal gives the “Index of Performance” (IP), also termed “bandwidth”, and measured in units of “bits per second”.

Although Fitts’ Law has been extensively studied for acquisition of on-screen targets, few studies have examined its effectiveness in modelling the acquisition of off-screen targets. Hinckley *et al* [2002] examined user performance with a variety of scrolling input devices (but no zooming), showing that Fitts’ Law accurately modelled off-screen target acquisition. Guiard *et al* [2004] describe two types of pointing involved in multi-scale (zoomable) off-screen target acquisition: *view-pointing* in which the user moves their view until the target is visible; and *cursor-pointing* in which the user moves the cursor over the final target. They theoretically examine the user’s movement through pan-zoom space using space-scale diagrams, predict conformance to Fitts’ Law, and empirically confirm the theory using bimanual parallel input for pointing (controlled by a stylus-tablet combination) and zooming (a joystick in the non-dominant hand).

3 Experimental interfaces

We developed three experimental interfaces from the same OpenGL/C++ core program: a text-document browser (Figure 1), a ‘flat’ map browser, and a globe map browser (Figure 2). The OpenGL graphics libraries allow rapid frame-rates and smooth animation through graphics hardware acceleration. Scrolling in all interfaces is controlled by rate-control, with the scroll-speed increasing linearly with the distance between the current and mouse-down cursor locations. Like the Microsoft Windows standard, all of our interfaces used the middle-mouse button to control rate-based scrolling. Each interface supported two zooming modes, either manual zooming controlled by the ‘a’ and ‘z’ keyboard keys, or automatic zooming in which the zoom-level is bound to the scrolling velocity. The relationship between mouse-displacement and scroll-speed was identical for each of the manual and automatic zooming pairs, as described in the subsections below.

The text-document browser only allows vertical scrolling, with the scrolling velocity controlled purely by vertical displacement of the mouse. The ‘flat’ map browser allows 2D scrolling in any direction up to the boundary of the map, with the scrolling velocity and direction dependent on the absolute distance between the mouse-down and current cursor locations. The globe browser also allows 2D scrolling in any direction, with the underlying globe rotating under the user’s cursor. In addition to the ‘global view’ of landmasses and oceans, the globe is populated with fifteen city maps that are represented as small coloured rectangles over each city’s location in the global view. When the user moves slowly or stops over a city the view rapidly zooms into the underlying map details.

In all interfaces the cursor is warped to the centre of the screen when the user begins scrolling. A red-arrow connects the screen-centre with the current cursor

location as the user drags the mouse. The arrow's direction shows the scrolling direction and the arrow's length depicts the scroll-speed.

The automatic zooming interfaces have a "maximum falling rate" which rapidly animates the transition between zoomed-out and zoomed-in views. Without a maximum fall rate there is a highly disconcerting effect of "slamming into the document" when the user stops scrolling by releasing the mouse button, or when reversing from rapid scrolling one direction to the other. All automatic zooming interfaces also supported a "scroll-to-cursor" function, which rapidly brings the document region under the user's cursor to the screen centre when they stop scrolling (by releasing the mouse button). Prior to implementing this function we found that users would often release the mouse button when the zoomed-out target was under their cursor, only to have the target fall outside the viewable region when the view returned to full-zoom. We observed that trial users' eyes typically followed the cursor when scrolling, and that they stopped scrolling when the cursor is over the target. The scroll-to-cursor function, therefore, brings the target to the screen centre through rapid animation. Issues associated with this function are discussed in Section 6.

The user's experience with zooming interfaces is strongly influenced by the precise calibration of the system's behaviour. This is particularly true of automatic zooming interfaces. To aid replication of our studies, exact details of the calibration settings for each of the interfaces (both manual and automatic zooming) are provided below, and they are summarised in Table 1. These values are based on theoretical and empirical analysis described by [Savage 2004] and in [Cockburn *et al* 2005].

3.1 Text-browsing interface

The text-browsing interface allows vertical document scrolling. Any Postscript or PDF document can be displayed, with the evaluations using a 157 Masters Thesis. On loading a document, each A4 page is converted into a 512×512 Targa Image File. Each page measured a true 21×27cm on the screen when rendered at 100% magnification on the displays used in the experiment.

Automatic and manual zooming both used a one-to-one relationship between vertical mouse displacement (in pixels) and resultant scroll-speed (in cm/second). Note that scroll-speeds are reported as document scroll-rates, rather than the rate that the pixels move across the screen—for example, at 50% magnification the document scroll-speed is twice the pixel movement rate.

In calibrating the behaviour of manual zooming, we were careful to make decisions that we felt would optimise its use. We were aware that leaving maximum velocities unconstrained at any particular zoom-level would allow pixel movement rates that exceed the capacity of the human smooth-pursuit visual system [Morgan and Benton 1989], yet we wanted to allow users to quickly accelerate to rapid document movement when zoomed out. Through informal experimentation with several trial users we decided to use five discrete zoom levels, activated through successive clicks of the 'a' (zoom in) and 'z' (zoom out) keys, each of which changed the magnification level by 17.5% between maximum and minimum zoom-levels of 100% and 12.5%. To reduce the disorienting effect

	Limits for automatic and manual zooming			Automatic zooming calibration of scroll-speeds for various magnification level						
	Max speed @ 100%	Min mag.	Max speed @ min mag	100%	75%	60%	50%	45%	25%	12.5%
Text (cm/sec)	48	12.5%	170	0-48	57	-	67	-	100	118-170
Map (cm/sec)	10	12.5%	80	0-10	20	-	30	-	40	50-80
Globe-view (deg/sec)	15	45%	60deg/sec	15	35	-	55	60	-	-
City-view (mins/sec)	5	60%	25	5	-	15-25				

Table 1. Speed and zoom calibration settings for the three interfaces with manual and automatic zooming.

of excessive scroll-speeds at each zoom-level, we applied maximum scrolling velocities at each zoom-level (in preliminary trials without the velocity caps several users complained of getting lost due to excessive speeds, particularly when ‘backing up’ after overshooting a target). Table 1 shows the maximum velocities at the 100% and 12.5% magnification levels. Maximum velocities for each of the four discrete magnification levels between 100% and 12.5% were determined by linear interpolation.

Calibration settings for the automatic zooming interface are also shown in Table 1. At speeds below 48cm/sec the document remains at full-zoom, but smooth zooming is applied beyond 48cm/sec through linear interpolation between the values shown. Between 118cm/sec and the maximum speed of 170cm/sec, the minimum zoom-level of 12.5% is applied.

3.2 Map-browsing interface

The map-browsing interface allows 2D rate-based scrolling over a detailed city map. The underlying map scrolls smoothly in whatever direction the user drags the mouse. The map used in the evaluation was a street map of Christchurch, New Zealand, displayed at 5120×3072 pixels (360×216cm on the screen at 100% magnification).

The mapping between mouse-displacement and scroll-speed was identical for both manual and automatic zooming modes, with a linear relationship of speed (cm/sec)=0.4×displacement (pixels), up to a maximum displacement of 200 pixels (and consequently 80cm/sec). The maximum map scroll-speed (80cm/sec) is lower than the maximum text document scroll-speed (170cm/sec) because the map continues to fill the display window at low zoom levels, while the text interface does not (Figure 1 shows that at low zoom levels the text window contains large blank regions).

Calibration of the manual zooming interface is similar to that of the document interface, with five discrete key-presses moving between full and minimum zoom levels of 100% and 12.5% respectively. Table 1 shows the maximum scroll-speeds

at the full and minimum manual zoom levels. Linear interpolation is used to determine maximum speeds at each discrete intermediate zoom level.

The automatic zooming calibration settings are also shown in Table 1. Linear interpolation is used to allow smooth zooming with changes in scroll-rate.

3.3 Globe-browsing interface

The globe-browsing interface is the most complex of the three due to its dual-view interaction of globe- and city-views. In the globe view, users can navigate around a globe representation of planet Earth, with rate-based scrolling causing the globe to rotate at an angular velocity proportional to the mouse-drag distance. Horizontal scrolling is unconstrained (the globe can endlessly rotate on its axis), but vertical scrolling is constrained to disallow rotation over the poles, without which the globe can be inverted, causing disorientation. Magnification levels in the globe view range from a minimum of 45%, showing the entire planet in a single window (see Figure 2), to 100%, which shows 20 degrees of arc at the equator in one window-width. The “100%” zoom level for the globe view was arbitrarily selected as the point at which further magnification of the images yields little benefit due to pixelation of the images.

The globe-browsing interface also supports a city view. Fifteen city maps are placed on the surface of the globe at a size of 2x1cm at 100% magnification on the displays used in the experiment. If the user slows or stops movement over one of the city maps, they zoom into a close view of the underlying map. The ‘fall’ into the city view is automatic when using automatic zooming, but is under explicit user control with manual zooming. Once in the city view, the system behaves similarly to the map interface (Section 3.2), except that dragging off the map ‘snaps’ back to the globe view with automatic zooming.

The relationship between mouse displacement and scroll-speed is different in the city and globe views. In the globe view there is a two-step relationship, allowing fine-control at slow speed and more coarse control at higher speeds: below 150 pixels displacement, scroll-speed (degrees/second)= $0.1 \times \text{displacement (pixels)}$; between 150 and 250 pixel displacement, scroll-speed= $-52.5 + 0.45 \times \text{displacement}$. In the city view, a linear relationship applies up to a maximum displacement of 200 pixels, with scroll-speed (minutes/second)= $0.1 \times \text{displacement}$.

With manual zooming there are five discrete zoom levels for the globe-view (between 100% and 45%) and a further five for the city-view (between 100% and 60%). Maximum scroll-speeds for the boundary conditions in both the globe and city views are shown in Table 1. Linear interpolation determines scroll-speeds for each of the intermediate zoom levels.

With automatic zooming, the zoom level is smoothly adapted to scroll-speed. Table 1 shows the relationship between speed and zoom in both the globe and city views. Linear interpolation is used between the values shown.

4 Experimental Details

The experiment is designed to answer two primary questions. First, does SDAZ (rate-based scrolling with automatic zooming) allow faster off-screen target acquisition than rate-based scrolling with manual zooming? Second, do users prefer SDAZ over rate-based scrolling with manual zooming, and do they find it less cognitively challenging? We also scrutinize the accuracy of Fitts' Law in modelling user performance in zoom-based off-screen target acquisition.

The participants' tasks involved acquiring a target depicted by a red rectangle that was 2x2cm on the screen when displayed at 100% magnification. The direction to the target was continually cued by a green arrow at the window centre. In the globe interface, the final target was always placed within a city view, and the city containing the target was highlighted red in the globe view. To complete the task, the user had to place the target, zoomed to 100%, under a cross in the window-centre and click the left mouse-button. Completing one task caused the next task to be generated, with the green arrow cueing the search direction. All user actions were continually logged by software.

It is important to note that these tasks do not require the user to extract and parse semantic information from the information space—rather, they mechanically 'chase' the red squares in the direction cued by the arrow, using rate-based scrolling plus either manual or automatic zooming. The decision to analyse mechanical interaction with the systems was intentional, as our prior work has already demonstrated that SDAZ better supports tasks that involve information seeking (although it did not investigate manual-zooming).

Thirty-five undergraduate Computer Science students (30 male, 5 female) took part in the experiment. All completed a questionnaire gathering background demographics regarding age, gender, dominant hand, and gaming experience. Training involved first watching a five minute demo of each of the three interfaces (text, map, and globe) in both zooming modes (automatic and manual zooming). They were explicitly instructed to combine scroll and zoom actions when using the manual-zooming interfaces. They were then given a few minutes to experiment with each zoom-type with each interface. Six practice tasks immediately proceeded each block of tasks with each interface and each zooming type. Data from the practice tasks were discarded. The order in which participants were exposed to each interface-type was controlled using a Latin square, and order of exposure to automatic and manual zooming was alternated across participants. Both manual and automatic zooming tasks were completed with each interface type (text, map, globe) before proceeding to the next interface type. NASA-TLX worksheets [Hart and Staveland 1988] were administered by software after each block of tasks, with participants using 5-point Likert scales to report various workload measures. On completing all tasks with each interface-type, the participants stated whether they preferred automatic or manual zooming.

4.1 Experimental design

The experiment is designed as a 2×6 repeated-measures analysis of variance (ANOVA) for factors zooming condition (automatic versus manual) and distance. The same experimental design is used to analyse data from each of the three interface-types: text, map, and globe. The factor ‘distance’ determines how far the target is placed from the starting position, as follows for the three interface types:

- Text interface—5, 10, 15, 20, 25, and 30 pages, where 1 page=27cm.
- Map interface—50, 100, 150, 200, 250, and 300 cm.
- Globe interface—4, 8, 16, 32, 64, and 128 degrees.

In the text and map interfaces, three tasks were completed at each distance with each zooming mode. In the globe interface two tasks were completed at each distance with each zooming mode. Therefore, including the six practice tasks, the tasks blocks consisted of 24 tasks with the text and map interfaces, and 18 tasks with the globe interface. Two tasks blocks were created for each interface, and the order that the blocks were used with each zooming condition was varied across participants.

To prevent extreme outliers, all acquisition times greater than the mean plus three standard deviations were removed from the analysis.

4.2 Apparatus

Participants used identical Athlon 1600+ computers with 256Mb of RAM running Linux 9.0, with Geforce 2 MX video cards outputting to 19-inch (36x27cm) Compaq monitors at 1280x1024pixel resolution. Input was provided through three-button Logitech mice with sample rates of 60Hz. The default RedHat 9.0 control-display gain settings were used: acceleration 2/1 pixels, threshold 4 pixels.

5 Results

Although the participants were able to complete almost all tasks quickly, they sometimes became ‘lost’ in the information space, overshooting the target and failing to attend to the green arrow directing them toward the target. As planned, we discarded ‘outlier’ tasks that exceeded the mean by more than three standard deviations. In total, 4.5%, 0.3%, and 0.8% of tasks were discarded with the text, map and globe interfaces respectively.

General observation of the participants indicated marked differences between the two zooming conditions, with much higher levels of concentration and physical activity when using manual zooming. These observations are supported by the analysis of NASA-TLX worksheets.

The following subsections present the analysis of variance of task completion times, then the Fitts’ Law modelling investigation, followed by the analysis of subjective measures of workload and preference.

5.1 Comparative performance analysis

Automatic zooming showed a small but statistically significant performance advantage over manual zooming with both the text and map interfaces (see Figures 3a & 3b). Mean task times with the automatic and manual conditions for text tasks were 6.8 (standard deviation 2.0) and 7.1 (s.d. 2.0) seconds respectively ($F_{1,34}=5.7$, $p<.05$), and for map tasks they were 5.6 (s.d. 1.7) and 6.2 (s.d. 2.0) seconds ($F_{1,34}=15.8$, $p<.01$). Mean task times for automatic and manual zooming with the globe interface (Figure 3c) were similar at 10.9 (s.d. 2.2) and 11.1 (s.d. 2.3), yielding no significant difference ($F_{1,34}=0.18$, $p=.7$).


As expected, there was a strongly reliable main effect for distance for all interfaces, but this simply confirms that tasks get harder as distance increases. More interestingly, while there was no interaction between factors zooming-condition and distance with the text and globe interfaces ($F_{5,170}<1$ and $F_{5,170}=1.9$, $p=.1$), there was an interaction with the map interface ($F_{5,170}=5.4$, $p<.05$). The cause of the interaction is visible in Figure 3b, which shows that performance with automatic zooming degrades less quickly than manual zooming as distance increases, particularly at high distances.

5.2 Fitts' Law analysis

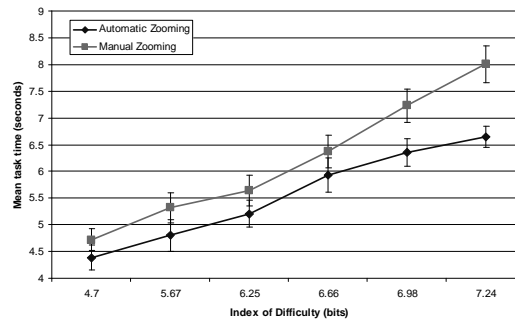
In the Fitts' Law analysis we used linear regression to calculate the line of best fit for the relationship between movement time and 'Index of Difficulty' (see Section 2.1). Index of Difficulty (ID), is calculated as $ID=\log_2(A/W+1)$, with W being the target size at 100% zoom (2x2cm or the equivalent number of minutes of arc in the globe-interface), and A being the total distance between the start and target locations at 100% zoom (in cms for text and maps, and in minutes of arc for the globe-viewer).

Fitts' Law accurately modelled user performance with both zooming types (automatic and manual) in the text and map interfaces. Table 2 shows the lines of best fit and the R^2 values for each interface-type and zooming condition. The good linear fits (with more than 90% of variance explained by the model) is normal for Fitts' Law pointing studies, and Hinckley *et al*'s investigation of scrolling to off-screen targets (without zooming) showed R^2 values in excess of 0.8.


Although data for the globe viewer initially suggests poor modelling by Fitts' Law, subsequent analysis reveals a good model ($R^2 > 0.8$) once the shortest distance tasks (4 degrees, or $ID=7$) are removed. The outlier poor performance with the short distance tasks is probably explained by a training effect from the majority of tasks: with the short-distance tasks, users would almost invariably overshoot their targets, snapping out of the city view and scrolling rapidly in the direction of the guiding arrow before realising they had overshoot the target.



(a) Text interface.



(b) Map interface.



(c) Globe interface.

Figure 3. Mean task times for the three interfaces with automatic and manual zooming across various index of difficulty values. Error bars show the mean \pm one standard error.

Interface-type	Zooming	Line of best fit	R ²	IP
Text	Automatic	MT=1.62*ID-5.82	0.92	0.62
	Manual	MT=1.64*ID-5.69	0.96	0.61
Map	Automatic	MT=0.92*ID-0.21	0.93	1.08
	Manual	MT=1.23*ID-1.44	0.82	0.82
Globe (all tasks)	Automatic	MT=0.22*ID+8.96	0.13	
	Manual	MT=0.60*ID+5.28	0.62	
Globe (7-bit tasks removed)	Automatic	MT=0.67*ID+4.11	0.84	1.49
	Manual	MT=0.98*ID+1.27	0.93	1.02

Table 2. Speed and zoom calibration settings for the three interfaces with manual and automatic zooming.

5.3 Subjective measures

The analyses above show a small performance advantage for automatic zooming over manual zooming. The subjective measures, however, reveal a large difference between the zooming conditions, confirming our informal observations that manual zooming demanded substantially more concentration and effort than automatic zooming.

The NASA-TLX worksheets divide workload into six categories: mental demand (concentration), physical demand (manipulation work), temporal demand (pace and time pressure), performance (self-sense of good performance), effort, and frustration. Table 3 summarises the results, measured from five-point Likert-scales, with ‘better’ interfaces producing lower values (low demands/effort/frustration or good performance). Automatic zooming uniformly received better mean scores, with all but two of the 18 metrics yielding significant differences (Wilcoxon matched-pairs tests). In particular, the high rating for “Physical demand” with the manual zooming interface (3.7, 3.6 and 3.9 for text, map and globe browsing versus 2.2, 2.1 and 2.6 with automatic zooming) supports our observations of heavy manipulation burdens arising from parallel bimanual input.

After completing all tasks with both zooming types for each interface type, the participants were asked to state which interface they preferred. In text tasks 27 preferred automatic zooming and 8 preferred manual ($\chi^2=9.3$, $p<.01$); in map tasks 26 preferred automatic zooming ($\chi^2=7.3$, $p<.01$); and in globe tasks 23 preferred automatic zooming versus 12 preferring manual ($\chi^2=2.9$, $p<.1$).

	Text Interface		Map Interface		Globe Interface	
	AZ	MZ	AZ	MZ	AZ	MZ
Mental Demand	2.3 (1.1)	2.9 (1.0)	2.3 (1.0)	2.9 (1.0)	2.7 (1.0)	3.2 (0.9)
Physical Demand	2.2 (0.8)	3.7 (0.8)	2.1 (0.9)	3.6 (1.0)	2.6 (0.9)	3.9 (0.9)
Temporal Demand	2.6 (0.9)	3.3 (0.9) →	2.5 (1.0)	2.9 (0.9)	2.6 (1.0)	3.0 (1.0)
Performance	2.0 (1.1)	2.3 (0.9)	2.0 (0.9)	2.5 (0.9)	2.1 (0.9)	2.6 (0.8)
Effort	2.6 (0.9)	3.3 (0.9)	2.5 (1.0)	3.1 (1.1)	2.9 (0.9)	3.4 (0.9)
Frustration	2.3 (1.0)	2.7 (1.2)	2.5 (1.1)	2.8 (1.0)	2.6 (1.0)	3.4 (0.8) →

Table 3. NASA-TLX Workload measures for Automatic and Manual Zooming (AZ and MZ) with the three interfaces. Mean (standard deviation) values shown, with lower values indicated lower workload or better performance. → indicates not significant at $p<.05$.

6 Discussion

The results show that automatic zooming allows faster off-screen target acquisition than manual zooming with less cognitive and manipulation effort. Automatic zooming was also strongly preferred to manual control. Finally, the results confirm that Fitts' Law is a robust model for zooming-based off-screen target acquisition.

In designing the experiment we were concerned to avoid experimental bias towards automatic zooming. We knew that by using Computer Science students as participants we were likely to have a high proportion of users who regularly play interactive computer games, which expose users to high rates of screen-based visual flow. We therefore collected background information on their gaming experience and used this to classify participants as 'gamers' or 'non-gamers' depending on whether they played interactive computer games for more than two hours per week. Through this scheme we divided our participant pool in sixteen gamers and seventeen non-gamers (nine of whom did not play games at all). We then compared gamer versus non-gamer performance with automatic and manual zooming in a 2x2 mixed factors ANOVA. Gamers outperformed non-gamers with the text interface (means of 6.2 versus 7.5 seconds, $F_{1,33}=19.5$, $p<.01$) and with the map interface (means of 5.2 versus 6.5, $F_{1,33}=10.1$, $p<.01$), but not with the globe interface (means of 9.8 versus 12.0, $F_{1,33}<1$). Importantly, however, there were no significant interactions between gamer-type and zooming type with any of the interfaces, meaning that there is no evidence to support the hypothesis that gamers are better able to exploit automatic zooming than non-gamers. This suggests that the benefits of automatic zooming should be available to a wide group of users.

There is also one reason for suspecting that our experimental method produced artificially favourable results for manual zooming. In order to generate the 'best-case-scenario' of performance with manual zooming we explicitly instructed users to use parallel bimanual controls by simultaneously zooming with the keyboard and scrolling with the mouse. Several commercial interfaces allow this style of interaction, but we suspect that most users overlook the capability because of the higher cognitive and manipulation workloads they demand. If our participants had not used parallel controls for zoom and scroll then their workload assessments for manual zooming would almost certainly improve, but at the cost of worse performance due to serial manipulation of scroll and zoom.

Another obvious experimental concern is that the tasks involved 'chasing a red blob' rather than meaningfully extracting information from the underlying information space. Again, this experimental design decision was intentionally made to better cover the sample-space of our previous and on-going research; while our previous work has focused on realistic information-extraction tasks, this experiment focuses on best-case mechanics of interaction.

Finally, several participants reported an important interaction problem with "hunting and overshooting" when using automatic zooming. One participant summarised the problem as "playing ping-pong over the final target". It seems that the problem was caused by the "scroll to cursor" function described in Section 3, which brings the portion of the information space under the cursor to the centre of the screen when the user releases the mouse button. While this technique works well at high velocities (that is, when zoomed out), it appears to work poorly when

scrolling slowly at full zoom. The participants seemed to naturally adapt to using scroll-to-cursor for approximate view pointing (“focus on the bit around here”), but they did not anticipate the same behaviour when scrolling at full-zoom. As a result, we believe that scroll-to-cursor should be disabled when scrolling at full-zoom.

7 Conclusions

Speed-dependent automatic zooming is an attractive interaction technique that automatically binds a document’s zoom level with its scroll-speed. In order to further test the effectiveness of automatic zooming, this paper investigated whether users can benefit from de-coupling the automatic relationship between speed and zoom, allowing users the freedom to explicitly control each property concurrently.

Results of a thirty-five participant study showed that participants completed tasks more quickly with automatic zooming than with manual zooming, that the user’s found automatic zooming less demanding and strongly preferable, and that the acquisition of off-screen targets is accurately modelled by Fitts’ robust model, even when scrolling is combined with zooming.

There is now substantial evidence that speed-dependent automatic zooming allows users to navigate through documents more quickly and with less effort than traditional document navigation techniques. In our further work we will conduct field studies of how our mature automatic zooming interfaces are used in everyday office work.

References

- Burr D, 1980 "Motion Smear" *Nature* **284** 164--165
- Casalta D, Guiard Y, Beaudouin-Lafon M, 1999 "Evaluating two-handed input techniques: rectangle editing and navigation", in *CHI '99 extended abstracts on human factors in computing systems* (Pittsburgh, Pennsylvania: Pittsburgh, Pennsylvania) pp 236-237
- Cockburn A, Savage J, 2003 "Comparing Speed-Dependent Automatic Zooming with Traditional Scroll, Pan and Zoom Methods", in *People and Computers XVII (Proceedings of the 2003 British Computer Society Conference on Human-Computer Interaction.)* Eds P Palanque, P Johnson and E O'Neill (Bath, England: Bath, England) pp 87-102
- Cockburn A, Savage J, Wallace A, 2005 "Tuning and Testing Scrolling Interfaces that Automatically Zoom", in *CHI'05. ACM Conference on Human Factors in Computing Systems* (Portland, Oregon: Portland, Oregon) pp 71-80
- Fitts P, 1954 "The Information Capacity of the Human Motor System in Controlling the Amplitude of Movement." *Journal of Experimental Psychology* **47** 381-391
- Guiard Y, Beaudouin-Lafon M, Bastin J, Pasveer D, Zhai S, 2004 "View Size and Pointing Difficulty in Multi-Scale Navigation", in *Proceedings of Advanced Visual Interfaces, AVI04* (Gallipoli, Italy: Gallipoli, Italy) pp 117-124

- Hart S, Staveland L, 1988 "Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research", in *Human Mental Workload* Ed P a M Hancock, N pp 139--183
- Hinckley K, Cutrell E, Bathiche S, Muss T, 2002 "Quantitative Analysis of Scrolling Techniques", in *Proceedings of CHI'2002 Conference on Human Factors in Computing Systems* (Minneapolis, Minnesota: Minneapolis, Minnesota) pp 65-72
- Hinckley K, Czerwinski M, Sinclair M, 1998 "Interaction and Modeling Techniques for Desktop Two-Handed Input", in *UIST'98. Proceedings of the ACM Conference on User Interface Software and Technology* pp 49-58
- Igarashi T, Hinckley K, 2000 "Speed-dependent Automatic Zooming for Browsing Large Documents", in *Proceedings of the 2000 ACM Conference on User Interface Software and Technology* (San Diego, California.: San Diego, California.) pp 139-148
- Leganchuk A, Zhai S, Buxton W, 1998 "Manual and Cognitive Benefits of Two-Handed Input: An Experimental Study" *ACM Transactions on Computer-Human Interaction* **5** 326--359
- MacKenzie I, 1992 "Movement Time Prediction in Human-Computer Interfaces", in *Proceedings of Graphics Interface '92*.
- Morgan M, Benton S, 1989 "Motion-deblurring in human vision" *Nature* **340** 385-386
- Savage J, 2004 *The Calibration and Evaluation of Speed-Dependent Automatic Zooming Interfaces* Masters Thesis, University of Canterbury, Christchurch
- van Wijk J, Nuij W, 2004 "A Model for Smooth Viewing and Navigation of Large 2D Information Spaces" *IEEE Transactions on Visualization and Computer Graphics* **10** 447-458
- Zhai S, Smith B, Selker T, 1997 "Improving Browsing Performance: A Study of Four Input Devices for Scrolling and Pointing Tasks", in *Proceedings of INTERACT'97: the sixth IFIP conference on Human Computer Interaction* pp 286-292