Spotlights: Attention-Optimised Highlights for Skim Reading

Byungjoo Lee
Aalto University, Finland
bjlee1985@gmail.com

Olli Savisaari
Aalto University, Finland
olli.savisaari@aalto.fi

Antti Oulasvirta
Aalto University, Finland
antti.oulasvirta@aalto.fi

ABSTRACT

The paper contributes a novel technique that can improve user performance in skim reading. Users typically use a continuous-rate-based scrolling technique to skim works such as longer Web pages, e-books, and PDF files. However, visual attention is compromised at higher scrolling rates because of motion blur and extraneous objects with overly brief exposure times. In response, we present Spotlights. It complements the regular continuous technique at high speeds (2–20 pages/s). We present a novel design rule informed by theories of the human visual system for dynamically selecting objects and placing them on transparent overlays on top of the viewer. This improves the quality of visual processing at high scrolling rates by 1) limiting the number of objects, 2) ensuring minimal processing time per object, and 3) keeping objects static to avoid motion blur and facilitate gaze deployment. Comprehension levels for long documents were comparable with those in continuous-rate-based scrolling, but Spotlights showed significantly better scrolling speed, gaze deployment, recall, lookup performance, and user-rated comprehension.

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation (e.g. HCI): User Interfaces – Interaction styles

Author Keywords

Skim reading; comprehension; scrolling techniques; attention brokering; visual attention; attentional blink

INTRODUCTION

This paper presents *Spotlights*, a novel technique to facilitate *skim reading* of long documents. The explosion of digital content on the Internet has made rapid comprehension of complex information important for users [24, 9]. Our goal is to improve the most commonly used interaction technique for documents, *continuous-rate-based scrolling*, which has become the *de facto* standard in browser and viewer applications. Scrolling rate is typically controlled by a scroll wheel, scrollbar, touchpad, or touchscreen.

The design of Spotlights addresses three factors that degrade information gain in continuous scrolling at high scrolling

Figure 1. Spotlights supports skim reading with continuous scrolling in document viewers. In scrolling down at a high scrolling rate, Spotlights selects important objects into semi-transparent overlays (spotlights) that stay static for a prolonged time before fading away. Thereby, the user can maintain visual attention on objects from previous pages.

speeds: 1) too many objects competing for visual selection, 2) overly brief per-object exposure times, and 3) motion blur [17, 30]. These drawbacks compromise users' ability to find information (lookup), explore contents, and comprehend documents when scrolling quickly. In particular, visual processing suffers if attention must be shifted too soon. This effect is called *attentional blink*. Current understanding is that about 500 ms of processing time is required to avoid it [29]. Although the phenomenon is more complex than this [26], the 'minimum 500 ms rule' provides a starting point for our approach. Spotlights is designed to complement continuous scrolling by giving a greater chance of spotting interesting objects at high scrolling speed. In our design, continuous scrolling takes place as normal below the rate of 2 pages/second (pps). Above this rate, Spotlights is triggered. It selects a few objects for static, semi-transparent layers (spotlights) rendered in the viewer window (see Figure 1). Their transparency allows the user to maintain a global view

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI'16, May 07–12, 2016, San Jose, CA, USA

Copyright © 2016 ACM. ISBN 978-1-4503-3362-7/16/05 ...\$15.00

DOI: <http://dx.doi.org/10.1145/2858036.2858299>

Figure 2. System overview: Spotlights extracts objects from a document in a series of preprocessing steps and renders them as an overlay in the renderer (viewer). The preprocessor saves only the meta-information of headings and figures to select overlay candidates from the image pool. The candidate objects can be highlighted as static overlays while the background image is scrolling.

of the scrolling document and spot potentially interesting objects underneath the layers. The spotlights remain on long enough to allow visual processing before fading away. Attention then is guided to another one appearing in the viewer. The detail-level behaviour of the technique is highly dynamic and responds to factors such as scrolling speed, clicking, and object type. Our approach requires no special hardware or expensive real-time computation. It can be integrated into a viewer that uses continuous scrolling. The only requirement is (lightweight) preprocessing of documents (see Figure 2).

Our technical contribution builds on and extends existing interaction techniques that select objects for attention at higher scrolling rates. We present a novel *design rule* to find the optimal spatial-temporal trade-off for human visual attention. More precisely, it determines 1) whether to show objects in spotlights / how many to show and 2) for how long. We claim that this supports visual attention in four ways: 1) divided attention is supported by limiting the number of objects for attention, 2) selective attention is supported by presenting the overlays in a predictable location, 3) motion blur is minimised by keeping overlays static (the gaze need not pursue them), and 4) attentional blink is avoided via an enforced minimum exposure time of 500 ms. The benefit to the user is that attention can be deployed to more objects and maintained longer at high scrolling rates.

To evaluate Spotlights, we followed a replicate-and-extend strategy wherein we repeated a set core task but implemented slight changes to conditions and dependent variables. To increase realism, we did not force users to use Spotlights but allowed them to control the scrolling rate in all studies. The three evaluation studies allow us to compare Spotlights to regular continuous-rate-based scrolling in terms of four central aspects of skim reading:

- Recall of contents (Study 1)
- Dynamics of scrolling (studies 1 and 2)
- Comprehension, self-reported (Study 1) and indicated by a questionnaire (Study 2)
- Lookup performance (Study 3)

The results are favourable. Spotlights not only supports keeping user attention on objects but improves interactions *during* scrolling. As users spot interesting items more reliably, they tend to do more backward scrolling. They are also faster in finding items they are looking for, and their recall of content improves. In other respects, scores were equal between Spotlights and the continuous-scrolling technique. We conclude the paper by discussing opportunities to exploit the properties of human visual attention in skim reading.

In the following discussion, we describe the design and implementation of Spotlights. We discuss how to handle issues such as masking and how to support page backtracking.

PREVIOUS WORK

We give a brief overview of skim reading, then position Spotlights in relation to existing scrolling techniques.

Definitions of skim reading describe it as being about *absorbing information* [11], wherein the defining feature is the aim of comprehending a new document in minimal time [24]. In addition to this main goal, skim reading involves sub-goals closer to the interaction. Firstly, it is an activity often involving *visual search tasks* wherein the goal is to find an object known in advance [24, 33, 7]. This emphasises *visual recognition* of objects during scrolling. Skim reading also involves *target acquisition*: after an interesting object is identified, it should be selected with the mouse or attended visually ([15]). Interestingly, empirical studies of unassisted skim reading suggest that most users are not efficient in it [3]. It has been noted that users exhibit a satisficing strategy wherein they jump ahead in the text when information gain falls below some threshold [9]. This is manifested in erratic scrolling rates.

Spotlights contributes to a body of techniques designed to support interaction with visual spaces that exceed the size of the viewport. These techniques typically target 1D visual spaces and documents. Five individual approaches can be distinguished. One is to make scrolling performance more robust for particular situations, such as dealing with varying lengths of documents ([6]) or single-handed mobile interaction ([25]). Secondly, the transition between pages can be enhanced. It has been found that users benefit from an animated transition more than page-flipping [22]. Rapid serial visual presentation (*RSVP*) is an alternative to page-based presentation that can minimise the motion blur in scrolling. It shows only static fragments of a document (e.g., words, phrases, or figures) at a time but at a very rapid pace. Thirdly, one can improve rate control, as with *hopping* [18] or *advance visual cues* [20]. A *wear-based scrollbar* [1] supports revisiting by visualising cues about previous visits to a page. Kim et al. [21] implemented a pseudo-haptic technique that creates an illusionary sensation of friction around important objects to alert users. Fourthly, one can relax the assumption that the document must be navigated linearly. Most 1D techniques here can be characterised via a scale-space diagram [10], including *Flipper*, which switches from continuous scrolling to flipping at 2 pps, and *Speed-Dependent Automatic Zooming (SDAZ)* [17] adaptively zooms out as scrolling rate increases. However, both suffer from decreasing resolution at higher

rates, making it hard to read headings and decipher figures. *Semantic rendering and zooming* have been proposed; these retain or enhance selected objects when the rate increases [19, 4, 27]. The fifth approach, which is orthogonal, is to train users to skim read more efficiently. Yi [31] presented *Qn-DReview*, a reading technique that facilitated rapid comprehension of scientific (CHI) papers.

Spotlights is a mono-scale technique that is best understood not in terms of scale but on the temporal dimension. As RSVP does, it selects objects for attention. Spotlights builds on RSVP by keeping these objects static, in overlays, but it mixes them with a rate-controlled background. As SDAZ's does, its behaviour changes with scrolling rate, but its selection and adaptation follow a novel design rule that prevents exceeding the limits of visual attention. This allows continuous scrolling at average rates of up to 10 pps, which pose a problem for existing techniques. Flipper, for example, prevents rates over 2 pps. Although our focus here is on the dynamic management of overlays, Spotlights is compatible with many of the above-mentioned techniques, including semantic zooming, advance cues, and hopping. One advantage of Spotlights is that no additional widgets or interactions are needed: it works in the familiar framework of a document viewer simply upon scrolling. Of the existing techniques, *Froggy GX* [32] comes closest. Aimed at enhancing Web page skimming for non-native learners of English, it uses content overlays, along with masking and filtering. However, it too lacks a solution to avoid overloading the visual attention.

OVERLAYS: DESIGN AND IMPLEMENTATION

The overlays in Spotlights are transparent, rendered in or near the position of the corresponding object on the original page, and fade out as a function of time. They are filled sequentially in first-come, first-served fashion, starting at the top of the display when multiple objects are shown simultaneously. We implemented the technique for documents composed of textual and graphical elements. However, the principles can be adapted for other types.

Extraction of Objects on a Page for Overlays

We automatically parse candidate objects for overlay that are likely to be interesting and descriptive. Previous work on skim reading suggests that users devote more gaze time to

Figure 3. Spotlights extracts objects for overlays by using horizontal Gaussian blur and blob detection. This identifies objects that are larger than their context.

Figure 4. Spotlights maximises performance by obtaining optimal exposure time E . Prolonged exposure allows more processing time for visual attention but causes overloading in the viewer window. Each object is guaranteed a minimum of 500 ms exposure, to avoid attentional blink.

titles, headings, pictures, and figures than plain text [24, 16, 9]. We used this fact as a guide in creating an object parser.

Our parser extracts headings and figures from a PDF page, using a two-step process. Firstly, the PDF file is converted to a set of images. Then, *horizontal* Gaussian blur is applied to each image. The result is a set of ‘smudged’ horizontal lines, as depicted in Figure 3. In the second step, we group these lines via a blob-detection process. We assume that there is a standard line spacing defined for a document, and those blobs of greater height than this line spacing are selected as possible important objects. Titles, headings, and figures are larger vertically than plain text and are likely to be selected in this process. Lastly, we sort the objects on a page in order of decreasing height, which allows us to select only the largest 1–2 objects per page for the overlays.

This parser is implemented in Java and works well for relatively homogenous content types (for other content, we may need a smarter extraction algorithm [5]). It reaches a processing speed of 100 pages per 40 seconds on a regular laptop. The technique is demonstrated in the supplementary video.

The Design Rule for Optimal Exposure Time E

After choosing objects, we render them in the location they originally had on the page. Here, we describe a design rule to decide how many objects to select and how long they should be displayed in overlays. The insight is that there is a trade-off between attention’s *spatial* resources – that is, space per object – and its *temporal* resources – that is, the time allowed per page. On one hand, it is beneficial if an object is shown for a long time. On the other, if exposure time is increased, more objects end up crowding the view. Therefore, exposure time should be increased only if there are few upcoming objects. Our design rule searches for the optimal exposure time for overlays, which maximises attentional performance while minimising other possible problems. We shall now describe the rule more formally.

Let us assume that a user is scrolling with N_{obj} important objects per page and a scrolling speed of v_p pps. From the visual attention perspective, this context affords $1/v_p$ seconds per page (T , the temporal resource) and $1/N_{obj}$ of space (or area) per object (S , the spatial resource). Spotlights affects T and S by choosing how long the objects remain there even after the page that contained them is no longer visible. Let us call this E , for exposure time of overlays. With longer E , more pages in the past leave objects in the current viewer space, hence

increasing the total number of objects per page. The changes in T and S due to exposure time E can be formulated in detail thus:

$$T = 1/v_p + E \text{ and } S = 1/(N_{obj}v_pE + N_{obj}) \quad (1)$$

However, if we multiply T and S , we always get the same value, no matter the changes in E , indicating temporal-spatial trade-off in the design space of Spotlights (see Figure 4):

$$T \times S = 1/(N_{obj}v_p) \quad (2)$$

Next, we explain the user perspective on the design rule. Let us say that the user is achieving a certain amount of attentional performance A with the given resources (T , S) and a given scrolling context (v_p , N_{obj}). Using Equation 1, we can express the function as shown below:

$$\text{Attentional performance} = A(T, S, v_p, N_{obj}) = A(E, v_p, N_{obj}) \quad (3)$$

If attentional performance is already at maximum and no bottleneck is occurring due to excess temporal and spatial resources, changes in E or other context parameters do not affect attentional performance, $A(E)$. In other words, the only bottleneck is in the user's attentional gate. However, here we suggest a design rule that first assumes a *temporal* bottleneck in the conventional scrolling context. It gradually increases E until we find that no more relief of the temporal bottleneck is available (see Algorithm 1). This 'staircasing' procedure is similar to just-noticeable difference (JND) methodology in psychophysics and was used in a previous study of layered user interfaces [12].

Algorithm 1 Design Rule to Find Optimal E

```

1: Define  $N_{obj}$ ,  $v_p$ ,  $\Delta E$ ,  $E_0$  and  $Threshold_E$ 
2: procedure SEARCH
3: top:
4:    $E_i \leftarrow E_{i-1} + \Delta E$ 
5:   Observe  $A_i = A(E_i)$ 
6:   if  $(A_i - A_{i-1}) < Threshold_E$  then
7:     return  $E_i$ 
8:   else
9:      $i \leftarrow i + 1$ 
10:  goto top.

```

In practice, the design rule can be parameterised in a user experiment. Firstly, N_{obj} , v_p , E_0 , and ΔE should be predefined on the basis of the scrolling context that we are striving to optimise. N_{obj} and v_p should not be extreme but represent the target range and type of documents. For E_0 , we would like to start at 500 ms to avoid attentional blink [28, 29]. Finally, the user's attentional performance should be appropriately observed. One of our studies used an eye tracker to measure attentional performance.

Dynamic Overlay Control

We adapt exposure time E for an overlay dynamically, in real time, in line with scrolling speed and user behaviour (see Figure 6). If the scrolling rate exceeds 2 pps, the overlay function is triggered. After triggering of the overlay, its transparency is linearly decreased until the set time expires. When transparency is high, it acts as a trigger for visual attention to move

Figure 5. When two or more overlays from the same location compete, they are 'stacked' in first-come, first-served fashion. An object's horizontal position is maintained.

on to the next object; also, it allows the attention to select the 'freshest' objects easily to focus on. The user can eliminate the overlays at any time by stopping the scrolling.

Our overlay control is content-specific. Taking into account the fact that textual information takes longer to comprehend [23], we obtain parameter E separately for graphical vs. textual objects, as described in the calibration section.

The overlay functions were implemented by means of an OpenGL framebuffer object. A frame rate of 60 Hz was achieved for 1050-pixel height with A4 or US-letter-sized documents.

Minimising Masking with Transparency

During scrolling, the movement of the document in the background layer is not meaningless; it gives the user a sense of scrolling speed and global position [22]. Hence, an overlay should not mask too much of the scrolling background. To minimise masking yet achieve high legibility of the overlay, we utilised an empirical result from a past study [12]. That finding was that a foreground opacity of 20% is required for obtaining good performance with both focused and divided attention. In our case, the overlay should get more focus than the background layer.

Stacking Objects

We also address competition for the same space among highlighted objects. This can occur when two overlays, from two or more previous pages, overlap in position. Titles and figures, for example, tend to be in the top half of a page. In this case, Spotlights tries to retain original positions, except that it treats the vertical and the horizontal dimension differently. Vertical positions for overlays are filled sequentially in first-come, first-served fashion, starting at the top of the display. In contrast, the horizontal position is kept as it was,

Figure 6. State diagram of the dynamic overlay control function in the Spotlights system.

Figure 7. Click-and-Go: The page containing the highlight can be warped to by clicking on it. This reduces the time for backtracking to spotted objects.

to retain some memory for object locations. When the document is crowded, the overlays are ‘stacked’. Figure 5 shows an example. The fact that the overlays are static and close to each other may aid in selective attention: a more predictable location may make it easier to find the next targets.

Click-and-Go: Page Warping with Selections

The overlays are selectable with the cursor. Before selection, the opacity control is turned off from the mouse-hover event (see Figure 6). Selection takes the viewer to the page containing the object in the overlay, in order to reduce the time for backtracking. Just after the warp, old highlights are removed from the viewer. Figure 7 illustrates this.

THE CALIBRATION STUDY

Here, we describe empirically obtaining exposure time E for textual and graphical objects. We applied our design rule (see Algorithm 1) for the following target scrolling scenario:

- $N_{obj} = 1\text{--}2$ objects per page
- $v_p = 10$ pps (averaged over 2–20 pps)
- $Threshold_E$ will be statistically tested with Helmer contrast
- $E_0 = 500$ ms, to avoid attentional blink
- $\Delta E = 250$ ms

Optimal exposure time E differs among various types of objects in a document. What matters is how quickly the user can recognise the object, absorb its ‘gist’, and attribute some value or importance to it. Our case is focused on two object types – headings (text) and graphical objects – but the procedure can be replicated for other types of object.

The experimental task was to follow rapidly scrolling symbols (shapes) or words. Users were asked to report which object variant (e.g., a green triangle) was the most frequent.

Figure 8. We obtained the optimal exposure time E for symbols (on the left) and textual objects (at right). Users were asked to attend a rapidly scrolling display to assess which of the three elements (e.g., rectangles, circles, and triangles) was the most common.

However, we did not examine this; we looked at only how long they needed to process an item before moving on to the next one. Figure 8 shows the task. As stimuli, we use simple symbols such as triangles and basic words. The reason is to obtain a *lower, optimistic boundary* for E , the minimal time for obtaining the gist of an object in rapid viewing.

Method

Participants: Four people (1 male, 3 female) were recruited from the local campus to participate. Their average age was 28.75 years ($SD = 2.98$), and three had corrected vision. All were familiar with scrolling.

Experimental design: The experiment followed a $2 \times 2 \times 6$ within-subjects design. The object-type conditions (2) and number-of-objects-per-page conditions (2) were balanced across participants. The duration of E was increased in six ascending levels, to avoid abrupt or surprising changes (this presumed that users get accustomed to the technique). We varied the type of object with two levels (word and symbol). The number of objects per page too was varied on two levels (1 and 2). The six levels for exposure time were 500, 750, 1000, 1250, 1500, and 1750 ms. The dependent measure of user performance was the time spent gazing at objects as a percentage of the trial’s duration (50 s).

Task and materials: We implemented a prototype of Spotlights using symbols (shapes) and words. We used a page size of 743×1050 pixels (a typical ratio for documents). Scrolling speed was fixed at 10 pps. Three variations of shape existed for the symbol condition: triangle, square, and circle. For the text mode, the name of the shape was displayed as a string (e.g., ‘triangle’ or ‘rectangle’; see Figure 8). The location and the colour of each object were randomly assigned when the page was generated. The object size was 50–150 pixels for a symbol. Font sizes between 25 and 50 pixels were used for text. Each trial presented 500 pages of objects generated randomly by the algorithm.

Procedure: Participants were asked to find a comfortable posture. The display was then set perpendicular to the field of view at a distance of 50 cm. After eye-tracker calibration, the experimenter explained the procedure. The participants were asked to attend to as many targets as possible during scrolling and at the end report which had been the most frequent shape during their scrolling. This task was devised to ensure that participants really attended to the stimuli. In reality, all shapes had the same probability. After the instruction, the experimenter started a trial consisting of scrolling through 500 pages. Because scroll speed was set to 10 pps, it took about 50 seconds to complete one trial. In total, it took about 30 minutes to complete the full 24 trials.

Apparatus: An SMI eye tracker (RED, 250 Hz) was installed just below an LCD monitor (Dell 21-inch with 60 Hz refresh rate; see Figure 8). A Java application on a laptop (MacBook Pro, late 2012, 15-inch, 2.2 GHz Intel Core i7, with Intel Iris Pro 1536 MB graphic card) generated and controlled the stimulus. Gaze location and the colour of the pixel at that location were saved in real time. We used colour information to extract the user’s attention to the highlighted objects.

Figure 9. At right, the optimal exposure time E for words vs. graphical objects. The estimates were obtained from eye-tracking data showing how long users were able to attend to highlighted layers as the number of objects and the exposure times changed. On the left, the average over the two object types.

Results

We used repeated-measures ANOVA with an alpha level of 0.05. Exposure duration showed a significant effect on the duration of attention to highlighted objects: $F(5,15) = 22.43$, $p = 0.001$. In particular, it increased with exposure time. However, the Helmert contrast shows no effect beyond Level 4 (1250 ms) of E ($p = 0.42$). Figure 9 presents this effect. The interaction effect between highlight duration and object type was also significant: $F(5,15) = 8.84$, $p < 0.001$. Pairwise comparisons prove no longer significant after Level 3 (1000 ms, $p = 0.11$) for symbols and Level 4 (1250 ms, $p = 0.08$) for words. This difference may stem from the differences in required withdrawal times between object types [23]. We conclude that words required 250 ms more exposure time.

The interaction effect between exposure duration and the number of objects on the page was not significant ($p = 0.16$). We conclude that at such a rapid rate, attention is already at its maximum with just a single object in an overlay.

STUDY 1: RECALL

Study 1 was informed by empirical research on skim reading, which emphasises the influence of 1) time pressure, 2) complex documents, and 3) focus on comprehension (as opposed to search). Our study imposed a time limit and used long documents (300+ pages). To assess comprehension, we asked users to recall keywords and figures, numerically rating their comprehension also. For directly testing whether Spotlights improves ability to attend to objects during scrolling, we also collected data on gaze and scrolling behaviour.

Method

Participants: Twelve individuals were recruited from the local university. Average participant age was 26.9 years ($SD = 6.44$). All had experience with scrolling. All reported being comfortable with the specific touchpad used for the study. Four users had corrected vision. The users were compensated with cinema tickets for participation. None reported having background knowledge of the materials.

Experimental design: The experiment followed a 2×2 mixed design with Scrolling Method and Book Set as factors. To mitigate order effects, we created two distinct sets of documents, and these were balanced between the two techniques. Each set consisted of three unrelated documents, and the

order of inner documents was counter-balanced. Scrolling Method was set as a within-subject factor while Book Set was set as a between-subject factor. Dependent measures included the number of figures and keywords recalled by the participants, participant-reported confidence in comprehension, the number of minutes the users would need to complete the skim reading, mean and maximum speed of scrolling, the amount of downward and backward scrolling, and depth of scrolling.

Task and materials: The participants were instructed to skim the whole document to get an overview. They could control scrolling speed as they wished. Participants were allowed to select their strategy for skim reading but asked to maximise global understanding of the document. Each trial was ended automatically after 90 seconds.

To increase realism, the materials used were books from various fields. The documents were chosen to have many pages ($M = 333$ pages, $SD = 195$). The books are listed below. Book Set 1 consisted of the first three books listed, and Book Set 2 was made up of the last three.

- *Architecture and Mathematics from Antiquity to the Future*, 723 pages (0.41 headings/page, 0.32 figures/page)
- *Higher Education in Sport in Europe: From Labour Market Demand to Training Supply*, 194 pages (0.98 headings/page, 0.29 figures/page)
- *Functional Programming in JavaScript*, 261 pages (0.55 headings/page, 0.019 figures/page)
- *Architectures: Modernism and After*, 270 pages (0.29 headings/page, 0.085 figures/page)
- *New Media, Old News: Journalism & Democracy in the Digital Age*, 233 pages (0.83 headings/page, 0.15 figures/page)
- *R Packages: Organize, Test, Document and Share Your Code*, 318 pages (0.52 headings/page, 0.054 figures/page)

These densities were computed by means of the object extractor of Spotlights. The average number of objects per page was 0.78 ($SD = 0.33$).

Spotlights was automatically triggered in the range of 2 to 20 pps. We used the parameters obtained in the calibration study. Because the average object density of the books used (0.75 objects/page) were below 1, we could use the empirically obtained E from the calibration study. Objects with a width-to-height ratio above 20 were regarded as textual objects and given 1500 ms of exposure time. Others were treated as figures and provided with 1250 ms. During the exposure period, the opacity of the background layer was decreased linearly from 1.0 to 0.2 on the basis of the instantaneous scrolling speed. For each page, the object with the greatest height was chosen as the candidate for Spotlights.

Procedure: The procedure was the same as in the calibration study, with the following exceptions. Participants controlled the scrolling with a touchpad on a laptop at a desk. The two-finger scrolling method (of Mac OS X) was used. The directional mapping was changed to suit the preferences of the user. After eye-tracker calibration, the experimenter demonstrated both scrolling methods in detail. Then, user was pro-

Figure 10. Study 1: Spotlights users reported higher confidence in their memory for objects ($p = 0.028$, left). Spotlights increased recall accuracy for graphical objects (right).

vided with five minutes for training with each method. A dummy document was used. After training, the experimenter explained the purpose of the study and the user's goal in it. Each trial lasted 90 seconds, and the Click-and-Go feature was turned off for this study.

Apparatus: We used the same set-up as in the calibration study.

Post-trial data collection: After each trial, four questions about understanding were asked. Two involved free recall and addressed memory for keywords and figures. The other two were about perceived comprehension. Participants were asked to rate their confidence in the level of comprehension they achieved on a 100-point scale. The quality of recall was not assessed, only the number of items.

Results

We used repeated-measures ANOVA with an alpha level of 0.05. Benjamini–Hochberg correction [2] was applied to correct for multiple testing, and all pairwise statistics used Bonferroni correction.

Free Recall

Memory of content improved with Spotlights. Scrolling method had a borderline effect on the number of figures recalled by the participants: $F(1,10) = 3.66$, $p = 0.11$. In particular, the mean for Spotlights was 2.69 ($SD = 1.38$) and for continuous scrolling 1.81 ($SD = 0.85$). However, no significant effect was found on the number of keywords recalled: $F(1,10) = 0.642$, $p = 0.50$. Overall, 7.73 keywords were reported ($SD = 2.81$). Scrolling Method $\times$ Book Set had a significant effect on the number of figures recalled: $F(1,10) = 6.914$, $p = 0.025$. Pairwise comparison showed that more figures were recalled ($p = 0.009$) when book set 1 was read with Spotlights ($M = 3.22$, $SD = 1.66$) than with normal scrolling ($M = 1.11$, $SD = 0.34$). For book set 2, there was no difference between scrolling methods.

Confidence

Users had more confidence in their comprehension with Spotlights: $F(1,10) = 6.61$, $p = 0.04$. In particular, Spotlights gained a 23.86% confidence score ($SD = 11.9$), while continuous scrolling received 15.33% ($SD = 11.2$).

No significant effect was found for the number of minutes users estimated they would need to complete the skim reading: $F(1,10) = 0.303$, $p = 0.59$. Participants reported, on av-

Figure 11. Study 1: Participants scrolled faster with Spotlights (left). Thanks to spotting more interesting objects, participants used more time scrolling backwards with Spotlights to revisit a passed object.

erage, that 10.79 minutes ($SD = 9.8$) would be required per document.

Pattern of Scrolling

Scrolling was faster with Spotlights. Mean scrolling speed differed significantly between the scrolling methods: $F(1,10) = 20.62$, $p = 0.009$. With Spotlights it was 3.26 pps ($SD = 1.53$), whereas continuous scrolling yielded 2.03 pps ($SD = 14.74$). Maximum scrolling speed too was significantly different: ($F(1,10) = 7.47$, $p = 0.047$). Spotlights showed maximum speeds of 35.19 pps, with continuous scrolling reaching only 25.64 pps.

Also, when using Spotlights, users scrolled significantly more in the downward direction: $F(1,10) = 9.97$, $p = 0.049$ – specifically, for 34.65 seconds ($SD = 12.74$) with Spotlights and 24.67 seconds ($SD = 16.96$) with continuous scrolling. However, the backward-scrolling time too was higher for Spotlights. Spotlights showed more backward scrolling ($M = 6.33$ s, $SD = 6.56$) than continuous scrolling did ($M = 2.04$ s, $SD = 2.08$) at rates exceeding 2 pps: $F(1,10) = 7.43$, $p = 0.038$. The effect of scrolling method on scrolling depth was also significant: ($F(1,10) = 9.83$, $p = 0.033$). Spotlights users scrolled through 209 pages ($SD = 68.3$), continuous-system users 157 pages ($SD = 79$).

Pattern of Gaze

We show eye-tracking data in Figure 12. The plot presents the movement speed of gaze vs. scroll speed. This figure shows

Figure 12. Study 1: Eye-tracking data show a more active gaze pattern for Spotlights. The amount of backward and forward scrolling in particular was more balanced with Spotlights (left), showing more vertical symmetry of shape than the normal scrolling condition did (right). In this plot, downward scrolling is expressed as negative scrolling speed.

that Spotlights allowed richer use of visual attention to objects, especially in fast scrolling. In other words, users could maintain visual attention longer on a larger number of objects in the document. Furthermore, the amount of backward and forward scrolling was more balanced with Spotlights, showing a vertically more symmetrical form than normal scrolling.

STUDY 2: COMPREHENSION

Our second study compared Spotlights to normal scrolling in a questionnaire-based comprehension test accompanied by standard workload metrics.

We learned from the recall study that participants employed more backward tracking. We wanted to assist in this behaviour by turning on the Click-and-Go feature (see Figure 7). When an object is highlighted, the user can select that item by clicking it with the cursor. This warps the user to that object in the document and removes any other highlighted objects from view. The user is able to return to potentially important and interesting objects within the document.

Method

Participants: Sixteen peoples were recruited from local universities. Most participants were women (3 male, 13 female), and ages ranged from 20 to 34, with an average of 25.6 ($SD = 4.3$). Six participants had corrected vision. Twelve of them assessed their English skills as high enough for comprehending documents. The other four reported their skills in English as at medium level. None had background knowledge of the testing materials. All participants were volunteers, and each was rewarded with one cinema ticket for this 60-minute study.

Experimental design: We followed a $2 \times 2 \times 2$ mixed design. Scroll Mode was varied on two levels (Spotlights and normal scrolling), as was Reading Time (2 minutes and 10 minutes). Also, two separate books were chosen, to create two distinct reading conditions:

- Condition 1: Reading an engineering book with Spotlights and an HCI book with normal scrolling
- Condition 2: Reading an engineering book with normal scrolling and an HCI book with Spotlights

Scroll Mode and Reading Time were set as within-subject factors, while Reading Condition was set as a between-subject factor. Order of conditions was fully balanced among participants.

Task and materials: The task was to skim read through two books, using the continuous-scrolling method while trying to comprehend as much information as possible. The task was carried out for two books: *Engineering: An Introduction for High School* and *Human-Computer Interaction (3rd Edition)*. Properties of the books are shown below:

- *Engineering: An Introduction for High School*, 220 pages (0.84 headings/page, 0.58 figures/page)
- *Human-Computer Interaction (3rd Edition)*, 255 pages (Chapter 1 to 7, 0.53 headings/page, 0.51 figures/page)

Participants were told to prepare for a post-reading questionnaire. Before proceeding to the actual experiment, participants used a demo application to gain familiarity with Spotlights.

Multiple-choice questionnaire for comprehension: The questionnaires were generated to assess skim reading comprehension, emphasising literal and reorganisational comprehension[8]. To simplify scoring, all questions were multiple-choice items with three alternatives, of which one was correct. This made the expected score one third the number of questions if a participant was guessing. The purpose was to measure comprehension of important aspects and content of the two books. The questions were devised to pertain to larger themes, body text, headers, figures, and graphs. No type of content was especially favoured or ignored. The goal in creation of the questions was to cover all important material and produce a test-like setup for measuring comprehension.

After all questions were ready, they were split into two distinct question sets per book (of 22 and 21 questions for the engineering book and of 18 questions each for the HCI book). This was done by separating them into beginning-, middle-, and end-of-book questions, then picking the same number of questions randomly from all three categories. Thus the entire book was represented in each of the question sets. Although the questions were picked randomly, their order was preserved such that it reflected the order in which topics were presented in the books.

The questionnaire was administered via an online form, which automatically recorded the responses. Some of the questions involved a picture to which a link was embedded in the questionnaire.

Procedure: The experiment was conducted in a noise-free room, with participants seated at a desk where they were asked to find a comfortable posture. A background questionnaire was presented to examine age, field of studies, level of English skills, etc. Participants were first told the general structure of the experiment and then instructed on how to use Spotlights. The experiment was carried out with a laptop. Each participant practised using the software for five minutes or as long as needed for feeling comfortable with it. Continuous scrolling was facilitated by the touchpad, with two-finger use.

After the practice, participants proceeded to the skim reading task. During the reading, a scrollbar indicating the current location of the viewing window was shown at the right edge of the document. The time left for reading was visualised also, at the left edge of the document, via a vertical progress bar. Participants were instructed to rest if there was any test time left after answering all questions. The experiment lasted one hour.

Apparatus: We used the same laptop as in Study 1, with the same Java application. The only difference was in the Click-and-Go function, which was turned on. The effective frame rate of the application was 60 Hz.

Figure 13. Time spent on reading, plotted against page number. Participants showed deeper exploration of documents when using Spotlights for the HCI book (right).

Post-trial data collection: After each reading task, the questionnaire was administered, related to the book. Question-set order was balanced among participants. Answering time was limited to eight minutes, which was sufficient for all participants. Absolute scores were divided by the total number of questions in the set to gauge comprehension. When the participant was done with one mode, the corresponding workload was measured with a NASA-TLX form [14] in addition to the questionnaire, and answering time was increased to 10 minutes. We used Raw TLX values [13] to simplify the experimental procedure.

Results

Repeated-measures ANOVA with an alpha level of 0.05 was used for statistical testing. All pairwise statistics used Bonferroni correction. One participant gave no answer to five questions on the engineering book and three on the HCI book. An ‘Incorrect’ response was recorded for these.

Comprehension Scores

Reading time had no significant effect on reading scores ($F(1,14) = 0.072$, $p = 0.79$). The score was 58% for two-minute reading ($SD = 9.4$), 61% ($SD = 10$) for the 10-minute condition. Scroll Mode had no significant effect on reading scores ($F(1,14) = 0.30$, $p = 0.60$). Spotlights showed an average score of 60% ($SD = 11$), and normal scrolling scored 61% ($SD = 7$). Reading Condition had no significant effect on reading scores ($F(1, 14) = 3.80$, $p = 0.071$). The average score with the engineering book was 61% ($SD = 9$); that for the HCI book was 59% ($SD = 11$).

Figure 13 visualises the reading time spent on each page of the books for all participants. We found that for books that were predictable to them (e.g., *Engineering for High School*), users tended to rely more on normal scrolling. For less familiar books (e.g., *Human-Computer Interaction*), users tended to move around the entire document freely, using Spotlights, until finding something interesting to investigate further.

To understand possible effects of the differences in reading pattern shown in Figure 13, we further divided all questions

into two groups: difficult (under 50% correct answers) and easy (above 50%). Of the 43 questions for the engineering book, 16 were regarded as difficult, and 11 out of 36 were regarded as difficult for the HCI one. For difficult questions, Scroll Mode $\times$ Reading Condition significantly affected comprehension score: $F(1,14) = 8.57$, $p = 0.01$. Most notably, Spotlights showed a 3.3% higher score ($p = 0.024$) than normal scrolling for reading condition 1. For reading condition 2, the difference was not significant. Scroll Mode $\times$ Reading Condition effect on score was not significant for easy questions: $F(1,14) = 0.38$, $p = 0.55$. We conclude that Spotlights helps users orient better to unfamiliar content, particularly by helping them understand the high-level structure of a book quickly.

Workload

None of the components in the TLX scale were significant: mental demand ($F(1,15) = 0.396$, $p = 0.54$), physical demand ($F(1,15) = 0.089$, $p = 0.77$), temporal demand ($F(1,15) = 0.016$, $p = 0.90$), performance ($F(1,15) = 0.791$, $p = 0.39$), effort ($F(1,15) = 0.127$, $p = 0.727$), and frustration ($F(1,15) = 0.281$, $p = 0.60$). Mean mental demand was 64.38 for Spotlights ($SD = 21.44$) and 66.56 ($SD = 20.63$) for normal scrolling. Physical demand with Spotlights was 27.19 ($SD = 25.03$), while that in normal scrolling was 25.63 ($SD = 23.16$); The result for temporal demand for Spotlights was 74.38 ($SD = 17.97$), 75 ($SD = 15.06$) for normal scrolling. For performance, Spotlights rated 65.31 ($SD = 20.37$), normal scrolling 61 ($SD = 21.49$). Effort was 68.44 for Spotlights ($SD = 17.49$) and 67.5 ($SD = 15.92$) for normal scrolling. Lastly, the corresponding figures for frustration were 52.81 ($SD = 23.59$) and 51.25 ($SD = 24.80$), respectively.

Usage Statistics for the Click-and-Go Feature

For the engineering book, participants used the Click-and-Go feature 7.06 ($SD = 7.19$) times with Spotlights. For the HCI book, they used this feature of Spotlights 11.56 ($SD = 9.63$) times.

STUDY 3: LOOKUP PERFORMANCE

The final study considered skim reading done for lookup and filtering purposes. A target was given and the user had to locate it in the document. Consider, for example, skimming the CHI proceedings to spot interesting articles for later reading. For such cases, scroll speed could be expected to be higher than in Study 2 (on comprehension), on account of the sparse distribution of interesting objects in the document. We again compare performance to continuous scrolling.

Participants: The same users who took part in Study 2 completed this study. Different materials and tasks were used, to avoid carry-over effects. Because of data loss caused by a technical error, two participants’ results had to be discarded, making the number of eligible participants 14. All were compensated with one cinema ticket for the 30-minute-long study.

Experimental design: Scroll Mode was the only factor in this within-subjects experiment. The conditions (Spotlights and normal scrolling) were balanced across participants.

Figure 14. Study 3: Lookup task screen for finding figures in 20 CHI full papers. A target figure was shown in the right portion of the screen, and the user's task was to find it by scrolling the document on the left. Spotlights users were twice as successful in localising targets as those using normal scrolling.

Task and materials: The document used in the task was a compilation of 40 CHI full papers from 2014 and 2015. They were randomly divided into two groups of papers thus:

- Group 1: 20 CHI full papers (200 pages, 0.38 figures/page)
- Group 2: 20 CHI full papers (198 pages, 0.65 figures/page)

The task was to find a given target figure by using the continuous-scrolling method. We limited the target type to figures because searches of the textual information could be done conventionally with the Find feature (Ctrl+F) of the PDF viewer. Thanks to their fixed similar formatting, CHI papers are good for our purposes.

A figure was randomly chosen from the entire range of documents and displayed at the right side of the viewer to specify the current target (see Figure 14). If a document was full of figures, as with an object density above 2 figures/page, Spotlights had to skip some figures to avoid overlap between overlays. However, such extreme cases were rare.

Procedure: The apparatus and the experimental setting were the same as for Study 2. The task here was to find as many figures as possible from the set of scientific articles. Each scrolling method was used for 10 minutes, with a short break between them, for a total task duration of just over 20 minutes. To indicate finding the figure, the Tab key was pressed on a keyboard. For a successful hit to be registered, the entire picture had to be inside the viewing window. If the key was pressed at any other time, a result of 'Incorrect' was recorded. At the beginning of the task for each method, the viewer showed the first page of the document. After each key press, a new picture was shown. The simplicity of the task meant that no practice session was needed.

The order of the pictures was randomised, and both mode order and document group were balanced between participants.

Results

We used repeated-measures ANOVA with an alpha level of 0.05 for the analysis. The scrolling method had a significant effect on the number of correctly found targets ($F(1,13) = 15.87, p = 0.002$). With Spotlights, users ($M = 15.5, SD = 7.76$) were twice as successful in localising targets as they were in normal scrolling ($M = 7.57,$

$SD = 3.08$). The number of false positives did not differ significantly between scrolling modes ($F(1,13) = 1.41, p = 0.26$). The mean of 'Incorrect' finds was 2.21 ($SD = 3.55$) for Spotlights and 1.00 ($SD = 1.03$) for normal scrolling.

SUMMARY AND FUTURE WORK

In this paper, we have presented a novel technique to support skim reading. It 'brokers' visual attention to objects in very fast scrolling of a document. As the user scrolls, it finds a balance between how many objects to show and how long to expose them. The benefits to users are threefold: 1) Even at high scrolling rates, some visual attention can be maintained and objects recognised. 2) The technique can be integrated into any application that uses continuous scrolling, on the assumption that (lightweight) preprocessing of documents is permissible. 3) It is intuitive and easy to learn. Our users learned it in five minutes during the study.

At a higher level, this paper has shown that we can approach the design of interaction techniques for skim reading in a systematic manner by considering the known limitations of visual attention. The results from our empirical study are positive and confirm that this approach deserves more attention. Specifically, we hypothesise that the attentional cost of shifting between overlays and background did not degrade performance through attentional blink, which prevents obtaining content-level information from underlying document at scrolling speeds over 2 pps. In summary, the results for Spotlights show that 1) users scroll 60% faster (3.26 pps vs. 2.03 pps) and scroll more; 2) they find more objects and track back more often to look at them; 3) they recall more figures, with a mild trend, and 4) are more confident in their level of comprehension; and 5) they succeed 100% better in lookup tasks when searching for given figures.

We see many opportunities to develop the technique further. Firstly, our exposure parameters were adjusted for a single scrolling speed. It is reasonable to assume that better performance would result if parameters were obtained for other rates too. Secondly, we observed that some users exhibit strategies that counteract our technique. Some users shifted attention to the middle of the screen instead of actively trying to find the best upcoming spotlight. It is useful to consider various techniques to guide attention almost automatically to the next object. The challenge is to do this in such a way that the document scrolling in the background is still legible. Thirdly, we did not consider how the complexity (or unpredictability) of an object affects visual processing requirements. Objects could be preprocessed for complexity and the exposure time modulated accordingly. Fourthly, we can imagine that models of individual-specific preferences or relevance could improve object selection. Fifthly, techniques such as advance cues or semantic zooming could be considered, to improve performance further.

ACKNOWLEDGEMENTS

The project has received funding from the Academy of Finland project COMPUTED and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 637991).

REFERENCES

1. Jason Alexander, Andy Cockburn, Stephen Fitchett, Carl Gutwin, and Saul Greenberg. 2009. Revisiting read wear: analysis, design, and evaluation of a footprints scrollbar. In *Proc. CHI09*. ACM, 1665–1674.
2. Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)* (1995), 289–300.
3. George Buchanan and Fernando Loizides. 2007. *Investigating document triage on paper and electronic media*. Springer.
4. George Buchanan and Tom Owen. 2008. Improving skim reading for document triage. In *Proc. IUI08*. ACM, 83–88.
5. Christopher Clark and Santosh Divvala. 2015. Looking Beyond Text: Extracting Figures, Tables and Captions from Computer Science Papers. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*.
6. Andy Cockburn, Philip Quinn, Carl Gutwin, and Stephen Fitchett. 2012. Improving Scrolling Devices with Document Length Dependent Gain. In *Proc. CHI12 (CHI '12)*. ACM, 267–276.
7. Andy Cockburn and Joshua Savage. 2004. Comparing speed-dependent automatic zooming with traditional scroll, pan and zoom methods. In *People and Computers XVII Designing for Society*. Springer, 87–102.
8. Richard R Day and Jeong-suk Park. 2005. Developing Reading Comprehension Questions. *Reading in a foreign language* 17, 1 (2005), 60–73.
9. Geoffrey B Duggan and Stephen J Payne. 2011. Skim reading by satisfying: evidence from eye tracking. In *Proc. CHI11*. ACM, 1141–1150.
10. George W Furnas and Benjamin B Bederson. 1995. Space-scale diagrams: Understanding multiscale interfaces. In *Proc. CHI15*. ACM Press/Addison-Wesley Publishing Co., 234–241.
11. Yves Guiard, Yangzhou Du, and Olivier Chapuis. 2007. Quantifying degree of goal directedness in document navigation: Application to the evaluation of the perspective-drag technique. In *Proc. CHI07*. ACM, 327–336.
12. Beverly L Harrison, Hiroshi Ishii, Kim J Vicente, and William AS Buxton. 1995. Transparent layered user interfaces: An evaluation of a display design to enhance focused and divided attention. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM Press/Addison-Wesley Publishing Co., 317–324.
13. Sandra G Hart. 2006. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, Vol. 50. Sage Publications, 904–908.
14. Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in psychology* 52 (1988), 139–183.
15. Ken Hinckley, Edward Cutrell, Steve Bathiche, and Tim Muss. 2002. Quantitative analysis of scrolling techniques. In *Proc. CHI02*. ACM, 65–72.
16. Jukka Hyönä, Robert F Lorch Jr, and Johanna K Kaakinen. 2002. Individual differences in reading to summarize expository text: Evidence from eye fixation patterns. *Journal of Educational Psychology* 94, 1 (2002), 44.
17. Takeo Igarashi and Ken Hinckley. 2000. Speed-dependent automatic zooming for browsing large documents. In *Proc. UIST00*. ACM, 139–148.
18. Pourang Irani, Carl Gutwin, and Xing Dong Yang. 2006. Improving selection of off-screen targets with hopping. In *Proc. CHI06*. ACM, 299–308.
19. Edward W. Ishak and Steven K. Feiner. 2006. Content-aware scrolling. In *Proc. UIST06*. ACM, New York, NY, USA, 155–158.
20. Victor Kaptelinin, Timo Mäntylä, and Jan ström. 2002. Transient visual cues for scrolling: an empirical study. In *Ext. Abst. CHI'02*. ACM, 620–621.
21. Juho Kim, Amy X Zhang, Jihee Kim, Robert C Miller, and Krzysztof Z Gajos. 2014. Content-aware kinetic scrolling for supporting web page navigation. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. ACM, 123–127.
22. Christian Klein and Benjamin B Bederson. 2005. Benefits of animated scrolling. In *CHI'05 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1965–1968.
23. Theresa J Babbitt Kline, Laura M Ghali, Donald W Kline, and Steven Brown. 1990. Visibility distance of highway signs among young, middle-aged, and older observers: Icons are better than text. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 32, 5 (1990), 609–619.
24. Fernando Loizides and George Buchanan. 2009. An Empirical Study of User Navigation during Document Triage. *Interfaces* (2009).
25. Takashi Miyaki and Jun Rekimoto. 2009. GraspZoom: zooming and scrolling control model for single-handed mobile interaction. In *Proc. MobileHCI09*. ACM, 11.
26. Mark R Nieuwenstein, Mary C Potter, and Jan Theeuwes. 2009. Unmasking the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance* 35, 1 (2009), 159.
27. Tom Owen, George Buchanan, Parisa Eslambochilar, and Fernando Loizides. 2010. *Supporting early document navigation with semantic zooming*. Springer.

28. Jane E Raymond, Kimron L Shapiro, and Karen M Arnell. 1992. Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance* 18, 3 (1992), 849.
29. Kimron L Shapiro, JE Raymond, and KM Arnell. 1997. The attentional blink. *Trends in cognitive sciences* 1, 8 (1997), 291–296.
30. Matteo Valsecchi, Karl R Gegenfurtner, and Alexander C Schütz. 2013. Saccadic and smooth-pursuit eye movements during reading of drifting texts. *Journal of vision* 13, 10 (2013), 8.
31. Ji Soo Yi. 2014. QnDReview: Read 100 CHI papers in 7 hours. In *Ext. Abst. CHI'14*. ACM, 805–814.
32. Chen-Hsiang Yu and Robert C. Miller. Enhancing web page skimmability. In *Ext. Abst. CHI'12*. ACM, 2655–2660.
33. Jian Zhao, R William Soukoreff, Xiangshi Ren, and Ravin Balakrishnan. 2014. A model of scrolling on touch-sensitive displays. *International Journal of Human-Computer Studies* 72, 12 (2014), 805–821.