AutoPager: Exploiting Change Blindness for Gaze-Assisted Reading

Andrew D. Wilson
Microsoft Research
Redmond, WA
awilson@microsoft.com

Shane Williams
Microsoft Research
Redmond, WA
shanewil@microsoft.com

ABSTRACT

A novel gaze-assisted reading technique uses the fact that in linear reading, the looking behavior of the reader is readily predicted. We introduce the AutoPager "page turning" technique, where the next bit of unread text is rendered in the periphery, ready to be read. This approach enables continuous gaze-assisted reading without requiring manual input to scroll: the reader merely saccades to the top of the page to read on. We demonstrate that when the new text is introduced with a gradual cross-fade effect, users are often unaware of the change: the user's impression is of reading the same page over and over again, yet the content changes. We present a user evaluation that compares AutoPager to previous gaze-assisted scrolling techniques. AutoPager may offer some advantages over previous gaze-assisted reading techniques, and is a rare example of exploiting "change blindness" in user interfaces.

CCS CONCEPTS

• **Human-centered computing → Interaction techniques**;

KEYWORDS

Eye tracking applications, gaze contingent displays, reading

ACM Reference Format:
Andrew D. Wilson and Shane Williams. 2018. AutoPager: Exploiting Change Blindness for Gaze-Assisted Reading. In *ETRA '18: 2018 Symposium on Eye Tracking Research and Applications, June 14–17, 2018, Warsaw, Poland*. ACM, New York, NY, USA, Article 4, 5 pages. https://doi.org/10.1145/3204493.3204556

1 INTRODUCTION

Today's eye trackers are unlikely to replace traditional pointing devices [Jacob 1991]. Meanwhile, *gaze contingent displays* more generally respond to the user's natural looking behavior [Baudisch et al. 2003; Duchowski and Çöltekin 2007]. A powerful example is *foveated rendering*, where graphics are rendered at the highest fidelity only in the region around the gaze point.

We propose the converse of foveated rendering: changing the graphics in the periphery, where the user *will* look. We demonstrate this approach in the context of linear reading (as when reading a novel), when looking behavior can be readily predicted. We introduce a "page turning" technique where the next bit of unread text is rendered in the periphery, ready to be read.

Gaze-assisted reading may be useful in hands-free interaction scenarios, such as when the user cannot use their hands or arms. Previous gaze-assisted reading techniques rely on scrolling text continuously or periodically in response to the user's gaze. In contrast, the proposed approach supports patterns of behavior more typical of reading novels in print or on e-reader devices. It does not require precise or fast eye tracking, nor high screen update rates.

2 RELATED WORK

Kumar and Winograd [2007] propose gaze-assisted text scrolling techniques, including continuous scrolling, where text smoothly scrolls to match the user's reading speed, and scrolling by page as the reader nears the bottom of the screen. Study participants found it disconcerting to read moving text, but became accustomed with practice. Participants preferred page-by-page scrolling, because it does not require reading moving text.

There have been few follow-on works to Kumar and Winograd. Räihä and Sharmin [2014] find that continuous scrolling results in smooth reading activity. Turner et al. [2015] propose a gaze-enhanced scrolling technique informed by observing when users manually scroll, supporting the user's preferred reading region.

This paper explores introducing new (unread) text into the periphery. This exploits the well-known "change blindness" phenomenon, where people are unaware of large changes clearly in view [Simons and Levin 1997]. Of particular relevance is the finding that when engaged in a cognitively and visually demanding task, viewers often miss extraordinary visual events [Santangelo et al. 2007; Simons and Levin 1997]. While there are detailed models of reading that describe how the reader perceives text in the neighborhood of where they are currently reading [Rayner 2009], we are aware of no work demonstrating change blindness while engaged in the cognitively demanding task of reading.

Some of the most powerful change blindness effects can be demonstrated when there is a visual disruption such as the display flashing black [Simons and Levin 1997]. However, change blindness can also occur in the absence of visual disruption, as when an object gradually fades into view [Simons et al. 2000].

In HCI change blindness is typically thought of as an undesirable effect [Davies and Beeharee 2012; Grad et al. 2007; Mancero et al. 2007]. Intille [2002] suggests that the potentially many displays in the user's periphery should be updated in a way that minimizes visual disruption. One proposed approach is to gradually fade in display updates. Dostal et al. [2013] demonstrate using an eye tracker to determine which of multiple displays the user is looking at, and

ETRA '18, June 14–17, 2018, Warsaw, Poland

ACM ISBN 978-1-4503-5706-7/18/06…$15.00
https://doi.org/10.1145/3204493.3204556

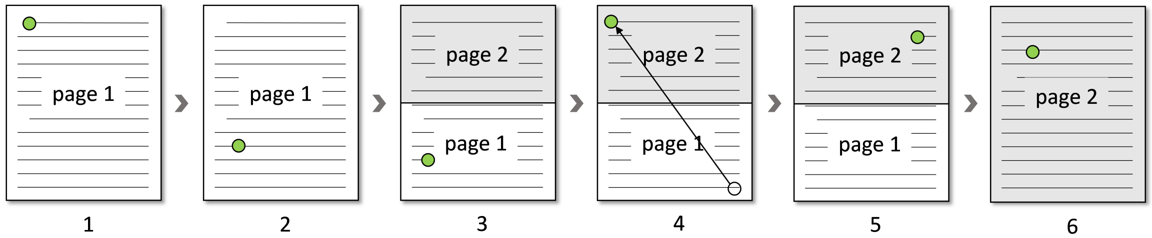

Figure 1: AutoPager automatically turns the page when linearly reading text (e.g., a novel). The green dot in the figures and in the accompanying video serves only to indicate where the user is looking, and is not present in the actual use of the system.

provide subtle rendering cues on the unattended display to indicate changes without drawing attention. Vidal et al. [2014] suggest introducing a notification in an augmented reality headset without attracting the viewer's eye by gradually increasing the brightness of the notification graphic only when the user is blinking, but this approach is not clearly demonstrated in their paper. Suma et al. [2011] and Azmandian et al. [2016] propose exploiting change blindness to reduce the discrepancy between the real and virtual in ways that are often go unnoticed in virtual reality.

3 AUTOPAGER TECHNIQUE

The proposed "AutoPager" technique uses an eye tracker to automatically turn the page when linearly reading text (e.g., a novel). The steps to turn a page below are illustrated in Figure 1:

1. The user begins reading from the top of the first page.
2. The user reads well into the bottom half of the page.
3. The top half of the second page is displayed; meanwhile, reading continues to the bottom half of the first page.
4. When finished reading the bottom half of the first page, the user saccades to the top of the page, where the top half of the second page is ready to read.
5. The user reads well into the top half of the second page.
6. The bottom half of the second page is displayed.

3.1 Transitions

The transitions from steps 2 to 3 and 5 to 6 may be designed to not distract the reader, and in the case of a slow cross-fade may be undetected [Simons et al. 2000]. In our prototyping we have empirically arrived at a gradual cross-fade effect (6s long) that minimizes the possibility of distracting the reader. Many readers are completely unaware of this subtle transition, and are left with the impression of reading the same page over and over again.

3.2 Advantages

The proposed technique has a number of features that set it apart from previous gaze-assisted reading techniques. Most importantly, the technique does not animate the text that the user is currently reading. The experience is more like that of reading a printed book or e-reader device. The saccade performed by the user to begin reading the next page is natural and reflexive. This contrasts with continuously scrolling techniques, where the reader may lack a sense of control or understanding of how scrolling works.

As our prototype implementation demonstrates by the use of two broad activation regions, the algorithm does not require an accurate or precise eye tracker. Other techniques that rely on spatial triggers or accurate measures of reading speed can behave erratically in the presence of eye tracker noise.

Because transitions are only coarsely synchronized with reading, the technique does not require high performance rendering, and therefore may be suitable for displays that update slowly (e.g., E-Ink displays), or eye trackers that exhibit high latency. Even in such cases, the next page's text will be available for reading well before the reader saccades from the bottom to the top of the page. Unlike with today's e-readers, there is no delay in turning the page.

3.3 Tracking the Reader's Progress

The technique requires that the system track the user's progress as they read from the top to the bottom of the page. For example, the transition of the top half of the page (steps 2 to 3 in Figure 1) should be performed only when the user has read well into the second half of the page, so that the transition affects regions of the display that are safely outside the foveal region of the reader, even in the presence of eye tracker noise. Meanwhile, the transition of the bottom half (steps 5 to 6 in Figure 1), should be performed only after leaving some time for the reader to possibly regress, or re-read, the last few words of the previous page.

Our prototype implementation divides the display into two activation regions, with the top region taking the top two thirds of the displayed page and the bottom region the bottom third (note that this differs from the even split as shown in Figure 1). The algorithm tracks how long the user has been reading in each region by incrementing a counter when the gaze point falls within the region [Starker and Bolt 1990]. To accommodate noisy eye tracking input that falls outside the display, these regions extend significantly beyond the display. The transition of the top half of the page is triggered after the reader has been looking at the bottom region for a total of 2.5s, while the transition of the bottom half of the page occurs after the reader has looked at the top region for a total of 5s, giving time to possibly regress to the bottom of the previous page. These threshold times were manually tuned for the page size, line spacing and reading speeds encountered during prototype development. Future implementations may benefit from determining these thresholds based on observed reading speed, etc.

One regression not addressed in the prototype is when the reader is reading in the bottom half of the page, the top half of the page has transitioned, and the reader then regresses to some point in the top half of the page (e.g., looking up to a previous paragraph). In this case, it may suffice to quickly undo the previous transition.

While short regressions and an immediate regression to the bottom of the previous page are supported, more non-linear reading such as skimming is not supported. Similarly, paging backwards and forwards may require an external mechanism, such page thumbnails or hardware buttons. To our knowledge all previously proposed gaze-assisted reading techniques have the same limitation.

3.4 Handling Interruptions

While reading text linearly, the reader may never detect that in steps 3-5 of Figure 1, the display consists of one half of one page and one half of the next. However, if the reader is interrupted and resumes reading at a later time, the illusion of reading coherent pages as if from a book may be broken. Worse, such an inconsistency may make finding their place in the text more difficult.

One strategy to handle interruptions is to force the display to render a single page, chosen based on where the reader left off. If the reader is reading the top half of the page before the interruption and the display is at step 5 in Figure 1, the system could force the transition to step 6. On the other hand, if the reader is reading the bottom half of the page and the display is at step 3, the system could move back to step 2. If upon returning to read, the reader resumes at a point near the top of the page, the system should move to step 4, with the bottom half of the previous page available for regression to the previous page. If the reader resumes at a point near the bottom of the page, the system should revert to step 3, with the top half of the next page available for reading.

If the reader resumes reading in the middle of the page, and they left off at a point nearby before the interruption, the system might revert to the step established before the interruption. If not, the course of action is less clear, since the user may want to regress to the previous page to re-establish reading location, or they may be ready to move on to the next page.

One approach to simplify resuming after an interruption is to provide the user with a visual indication of where they left off (akin to Kumar and Winograd's "gaze marker"). This indicator would be removed once reading is resumed.

4 EVALUATION

We conducted a user study comparing AutoPager with two previous automatic scrolling techniques. Twelve adult participants (four female) were recruited from the authors' institution. All had normal or corrected vision, were proficient in reading English, and had not been exposed to the techniques used in the study. Participants were compensated $10 for the study, which took less than one hour.

4.1 Apparatus

Participants sat in an office chair that had no wheels and could not swivel. On the desk in front of them was placed a Dell S2409W display (24″ diagonal, 1920x1080, 60Hz refresh), a mouse and keyboard. A SensoMotoric Instruments (SMI) REDn eye tracker [SensoMotoric Instruments, GmbH 2016] was mounted to the bottom bezel of the display in the manner directed by the manufacturer. User's distance to the eye tracker-equipped display was approximately 500-700mm, within the "eye-box" of the eye tracker. According to the manufacturer's specification, the SMI REDn eye tracker delivers 0.4° gaze position accuracy. All participants completed the manufacturer-supplied calibration procedure.

4.2 Task

The study consisted of multiple sessions of reading onscreen text, one session per condition. Text was sourced from the Wikipedia entry for Stanley Kubrick's film *2001: A Space Odyssey*, reduced to plain text by removing all graphics, links and other formatting. This was divided into several short contiguous sections, each read to completion and in the order in which they appear in the original article, regardless of the order in which conditions were presented (i.e.; all sections are read under all conditions).

Black text was rendered on a 160mm wide by 187mm high white rectangle centered on the display, with 20 lines of 22 point Times New Roman font. This emulates popular e-readers such as the Kindle, and gives a good number of pages to be read in a short time.

Before beginning the reading session trials, participants were notified that the experimenter would be observing their gaze patterns closely using a second display showing the text and the gaze point, and that it was expected that they read the text completely without breaks. Our informal observations of the second display confirm that participants complied with instructions.

We implemented two comparison techniques, drawing from Kumar and Winograd [2007]. We were primarily interested in determining if our proposed technique was usable, and whether it offered any advantages or disadvantages over previous techniques. We tested the following conditions within-subject:

Continuous Scrolling: Reading begins at the top of the page. When reading passes a fixed threshold line halfway down the screen, the page begins to smoothly scroll upwards, revealing new text at the bottom. Scrolling rate is matched to reading speed, such that reading occurs mainly around the same position onscreen. The threshold line was manually tuned to be about 60% down the page.

Page Scrolling: When reading reaches the bottom 30% of the page, the page smoothly scrolls 12 lines up (60% of the page height) over 1.5s. During the animation the area around the last reading position is highlighted by dimming the page outside this region. Highlighting is removed when the scrolling animation is completed. This effect, roughly analogous to Kumar and Winograd's "gaze marker", assists the reader in recovering their reading position after the page is scrolled, but does not prevent reading of text in the highlighted region during the animation.

No Scrolling: The AutoPager technique with a 6s cross-fade. So as to not tip off participants to the test condition, the technique was never referred to by its name.

4.3 Procedure

Instructions to participants noted only that various gaze-assisted reading techniques were to be tested, and that they used an eye tracker. Presentation order was counterbalanced by presenting each order of three conditions twice. Participants first read a section of practice text with one of the three reading techniques, and then read a section of the test text to completion using the same technique. Each section of text was about 5,000 characters long.

4.4 Results

While some people clearly read faster than others, a large difference in reading speed between techniques, aggregated over all participants, may indicate problems in a technique's usability. For example, a reader may pause, slowdown, regress (re-read), or even lose their place when distracted by an animation or transition.

Individual differences in reading speed were large across all test reading sessions (min. 149 WPM, max. 533 WPM), yet there was little difference in reading speed between conditions, with average reading speed 343 (s.d. 113) WPM, 334 (s.d. 116) WPM and 328 (s.d. 109) WPM for *Continuous Scrolling*, *Page Scrolling* and *No Scrolling* conditions, respectively. Analysis of variance of reading speed should illuminate differences between techniques while taking into account individual differences in reading speed. Repeated measures ANOVA yielded no statistically significant differences between reading technique conditions ($F_{2,11} = 2.29$, $p = .12$). We cannot conclude one technique is more effective than others based on differences in reading speed. This is consistent with our informal observations of the participants' eye gaze point on the secondary display; we saw no serious disruptions or obvious slowdowns.

In the *No Scrolling* condition, we saw only a few instances of page regression, when the reader briefly refers to the bottom of the previous page before continuing to read the top of the next page. In the practice sessions, there were a few instances of the participant looking away from their current reading place to investigate the transition, and none during test sessions.

In the *Page Scrolling* condition, many participants at first stopped reading during the scrolling animation. With practice, some participants began to read during the animation, albeit somewhat slowly.

As indicated by reading speed and informal observations, participants' reading performance was similar across all conditions, leaving little opportunity for further quantitative analysis.

Survey questions were designed to call out differences between the techniques, particularly as they relate to how the text is modified onscreen and the participants' understanding of how the techniques work. We performed repeated measures ANOVA on Likert scale questionnaire responses (1 = disagree, 7 = agree). Means and standard deviations of questions with ratings showing statistically significant differences across conditions are shown in Figure 2. We summarize some of the more interesting survey results:

"I was aware of when new text is introduced to the display": Responses are consistent with informal observations that the gradual transition in the *No Scrolling* condition was often undetected.

"I could read as slow as I like": Responses indicate that when using the *Continuous Scrolling* technique, people felt somewhat out

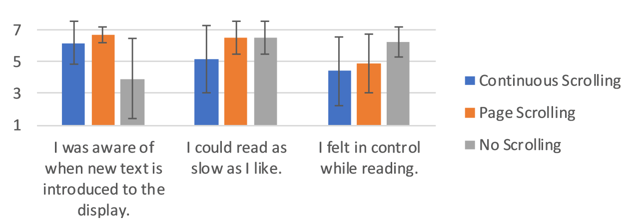

Figure 2: Means and standard deviations of questions with ratings showing statistically significant differences ($p < .05$).

of control as text continued to scroll. Some participants commented that they felt as if they were "on a conveyer belt".

"I felt in control while reading": Responses are consistent with the observation that in the *No Scrolling* condition, the reader takes no apparent, explicit action to advance the text while reading.

Participants were also asked to rank the techniques in order of preference. Six participants ranked *No Scrolling* as their favorite technique. Four ranked *Continuous Scrolling*, and two ranked *Page Scrolling* as their favorite technique.

Free form written comments collected after each condition show some themes. Regarding the *No Scrolling* technique, one participant noted that "I found *No Scrolling* to have the biggest wow factor. I enjoyed that the new text was simply there when I was ready for it without me doing anything (the other two still required me to follow the current line as it moved up the page)". Another commented that "*No Scrolling* felt very natural. If this technique could switch back to the top of the page if I didn't complete the entire page, that would be perfect."; this participant never attempted a regression and did not know that this feature was already supported. Another commented that it was "more mysterious than natural on first use".

Seven participants noted the lack of support for non-linear reading in the *No Scrolling* condition. One noted that "it might be appropriate for reading a story, however I am not sure how I could go back or skip pages if this was a technical document or a study document."

Several participants felt pushed to read faster in the *Continuous Scrolling* condition. One "found it to be tiring on the eyes; felt out of control". Another commented: "I could feel it slow down or speed up as my reading speed changed; somehow I felt pressured to read faster when it slows down". Another favored the technique: "it requires little effort from my part, and there is no interruption while reading". Four participants noted the desire to stop reading during the scrolling animation of the *Page Scrolling* condition.

All three reading techniques appeared usable. Many participants commented on the "wow" factor of the *No Scrolling* conditions, and half favored this condition. Interestingly, while some participants noted the lack of the ability to skim or quickly turn pages, they did not comment on similar limitations of the comparison techniques.

5 EXTENSIONS AND CONCLUSION

The present study does not analyze the impact of eye tracker precision, accuracy, and latency. In fact, the basic technique may require minimal eye tracking performance since it only requires determining which half of the display is being looked at, and is tolerant to latency. It may be possible to use simple, coarse pupil tracking-based techniques on front-facing mobile cameras, for example.

The AutoPager technique might apply to text formats other than a single column. For example, in addition to updating the top and bottom halves of a single column page, it may be possible to update the left and right columns of a two column format paper, or even the left and right halves of a small single line display.

While the Autopager technique addresses linear reading only, we believe there is opportunity in exploiting change blindness in user interfaces more generally.

REFERENCES

Mahdi Azmandian, Mark Hancock, Hrvoje Benko, Eyal Ofek, and Andrew D. Wilson. 2016. Haptic Retargeting: Dynamic Repurposing of Passive Haptics for Enhanced Virtual Reality Experiences. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1968–1979. https://doi.org/10.1145/2858036.2858226

Patrick Baudisch, Doug DeCarlo, Andrew T. Duchowski, and Wilson S. Geisler. 2003. Focusing on the Essential: Considering Attention in Display Design. *Commun. ACM* 46, 3 (March 2003), 60–66. https://doi.org/10.1145/636772.636799

Thomas Davies and Ashweeni Beeharee. 2012. The Case of the Missed Icon: Change Blindness on Mobile Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 1451–1460. https://doi.org/10.1145/2207676.2208606

Jakub Dostal, Per Ola Kristensson, and Aaron Quigley. 2013. Subtle Gaze-dependent Techniques for Visualising Display Changes in Multi-display Environments. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces (IUI '13)*. ACM, New York, NY, USA, 137–148. https://doi.org/10.1145/2449396.2449416

Andrew T. Duchowski and Arzu Çöltekin. 2007. Foveated Gaze-contingent Displays for Peripheral LOD Management, 3D Visualization, and Stereo Imaging. *ACM Trans. Multimedia Comput. Commun. Appl.* 3, 4, Article 6 (Dec. 2007), 18 pages. https://doi.org/10.1145/1314303.1314309

Kevin Grad, T. C. Nicholas Graham, and A. James Stewart. 2007. Effective Use of the Periphery in Game Displays. In *Proceedings of the 2007 Conference on Future Play (Future Play '07)*. ACM, New York, NY, USA, 69–76. https://doi.org/10.1145/1328202.1328215

Stephen S. Intille. 2002. Change Blind Information Display for Ubiquitous Computing Environments. In *Proceedings of the 4th International Conference on Ubiquitous Computing (UbiComp '02)*. Springer-Verlag, London, UK, UK, 91–106. http://dl.acm.org/citation.cfm?id=647988.741489

Robert J. K. Jacob. 1991. The Use of Eye Movements in Human-computer Interaction Techniques: What You Look at is What You Get. *ACM Trans. Inf. Syst.* 9, 2 (April 1991), 152–169. https://doi.org/10.1145/123078.128728

Manu Kumar and Terry Winograd. 2007. Gaze-enhanced Scrolling Techniques. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology (UIST '07)*. ACM, New York, NY, USA, 213–216. https://doi.org/10.1145/1294211.1294249

Gabriela Mancero, William Wong, and Paola Amaldi. 2007. Looking but Not Seeing: Implications for HCI. In *Proceedings of the 14th European Conference on Cognitive Ergonomics: Invent! Explore! (ECCE '07)*. ACM, New York, NY, USA, 167–174. https://doi.org/10.1145/1362550.1362584

Kari-Jouko Räihä and Selina Sharmin. 2014. Gaze-contingent Scrolling and Reading Patterns. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational (NordiCHI '14)*. ACM, New York, NY, USA, 65–68. https://doi.org/10.1145/2639189.2639242

Keith Rayner. 2009. Eye movements in reading: Models and data. *Journal of eye movement research* 2, 5 (2009), 1.

Valerio Santangelo, Marta Olivetti Belardinelli, and Charles Spence. 2007. The suppression of reflexive visual and auditory orienting when attention is otherwise engaged. *Journal of Experimental Psychology: Human Perception and Performance* 33, 1 (2007), 137.

SensoMotoric Instruments, GmbH 2016. Introduction to Bayesian Statistics. (2016). Retrieved April 12, 2016 from http://smivision.com/en/gaze-and-eye-tracking-systems/products/redn-scientific.html

Daniel J Simons, Steven L Franconeri, and Rebecca L Reimer. 2000. Change blindness in the absence of a visual disruption. *Perception* 29, 10 (2000), 1143–1154.

Daniel J Simons and Daniel T Levin. 1997. Change blindness. *Trends in cognitive sciences* 1, 7 (1997), 261–267.

India Starker and Richard A. Bolt. 1990. A Gaze-responsive Self-disclosing Display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '90)*. ACM, New York, NY, USA, 3–10. https://doi.org/10.1145/97243.97245

Evan A Suma, Seth Clark, David Krum, Samantha Finkelstein, Mark Bolas, and Zachary Warte. 2011. Leveraging change blindness for redirection in virtual environments. In *Virtual Reality Conference (VR), 2011 IEEE*. IEEE, 159–166.

Jayson Turner, Shamsi Iqbal, and Susan Dumais. 2015. Understanding Gaze and Scrolling Strategies in Text Consumption Tasks. In *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers (UbiComp/ISWC'15 Adjunct)*. ACM, New York, NY, USA, 829–838. https://doi.org/10.1145/2800835.2804331

Mélodie Vidal, David H. Nguyen, and Kent Lyons. 2014. Looking at or Through?: Using Eye Tracking to Infer Attention Location for Wearable Transparent Displays. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers (ISWC '14)*. ACM, New York, NY, USA, 87–90. https://doi.org/10.1145/2634317.2634344