Exploring Pointer Assisted Reading (PAR): Using Mouse Movements to Analyze Web Users' Reading Behaviors and Patterns

Ilan Kirsh[1](✉) and Mike Joy[2]

[1] The Academic College of Tel Aviv-Yaffo, Tel Aviv, Israel
kirsh@mta.ac.il
[2] University of Warwick, Coventry, UK
M.S.Joy@warwick.ac.uk

Abstract. This paper explores Pointer Assisted Reading (PAR), a reading behavior consisting of moving the mouse cursor (also known as the pointer) along sentences to mark the reading position, similarly to finger-pointing when reading a book. The study shows that PAR is an uncommon reading technique and examines methods to extract and visualize the PAR activity of web users. An analysis shows that PAR data of real users reveal reading properties, such as speed, and reading patterns, such as skipping and rereading. Eye-tracking is usually used to analyze user reading behaviors. This paper advocates for considering PAR-tracking as a feasible alternative to eye-tracking on websites, as tracking the eye gaze of ordinary web users is usually impractical. PAR data might help in spotting quality issues in the textual content of a website, such as unclear text or content that might not interest the website users, based on analyzing reading properties and patterns (e.g. reading speed, skipping, and rereading). Accordingly, PAR-tracking may have various practical applications in a wide range of fields, and particularly in educational technology, e-learning, and web analytics.

Keywords: Mouse pointer · Mouse cursor · Eye-tracking · Website · Web pages · Reading · Human-computer interaction · Educational technology · E-learning · Distance learning · Web analytics · Document · Text · Sentence · Word

1 Introduction

Reading habits and behaviors vary from person to person. This paper explores the behavior of reading text online with the aid of a mouse. Moving the cursor (also known as the pointer) of a pointing device (e.g. a mouse or a touchpad) to point at sentences and words while reading can be referred to as Pointer Assisted Reading (PAR) [16].

C. Stephanidis et al. (Eds.): HCII 2020, LNCS 12424, pp. 156–173, 2020.
https://doi.org/10.1007/978-3-030-60117-1_12

Previous studies have already shown proximity between the mouse cursor position and the user's eye gaze on the screen during mouse activity (as discussed in Sect. 2). Based on this knowledge, we can consider mouse cursor movements along lines of text in the direction of reading (left to right in English) as representing eye gaze movements, and accordingly, reflecting the reading of these lines of text.

The primary purpose of this study is to explore and examine PAR activity and to investigate and illustrate reading behaviors that are reflected in PAR data, such as changing reading speed, sentence skipping, and word rereading. This study paves the way to exciting, new research directions and practical applications.

This paper is organized as follows. Section 2 presents related work. Section 3 introduces the implementation of tracking and visualizing mouse movements, which was used in this research to collect, manage, and present relevant data. Section 4 shows that PAR is uncommon, proposes methods for "mining" PAR activity from mouse movement data, and presents experimental results that demonstrate the effectiveness of these methods. Section 5 presents and analyzes examples of PAR sessions of real users. Section 6 discusses the results. Lastly, Sect. 7 concludes this paper and suggests possible further work.

2 Related Work

User attention data are essential in many applications, including in educational technology, online learning, e-commerce, news websites, online advertising, and web analytics, as well as in studying reading habits and behaviors.

Eye and gaze tracking data have been used for estimating user attention in many applications and studies, including, for example, to assess the mental workload that a user interface places on users and to identify usability issues [13], to measure levels of attention to advertising [22], to guarantee user attention on application permission authorization [15], to measure mobile web user interest [23], and to evaluate the effect of social information on user enjoyment of online videos [25].

User attention data are particularly important in analyzing reading behaviors. Eye-tracking has been used in many studies on reading in recent years, including to distinguish between reading and skimming [2], to explore second language vocabulary acquisition while reading [27], to examine the effect of location-driven logo placement on attention and memory [11], to assess the effects of listening to music while reading [34], to explore the reading development of elementary school children over time [31], to identify reading disabilities and dyslexia [1], and to compare first-language reading to second-language reading [7].

Although eye-tracking technology is powerful and accurate, it has significant availability and scalability limitations. Technically, integrating eye-tracking into web pages is possible [8], but it is impractical for most websites, as eye-tracking normally requires special equipment on the client-side to capture accurate data, it requires user collaboration, and it can raise privacy concerns, as it makes use of cameras. A common solution is to use the alternative client-side user actions, such as page scrolling, mouse movements, and clicks, as implicit indicators of user attention [5,18,33]. Tracking user actions in modern browsers can be achieved by embedding special client-side JavaScript code in web pages. At least 15 commercial web analytics services track user mouse activity [17].

By collecting data on the client-side, the evaluation of user attention at any point in time can be easily reduced from the entire web page to the viewport, which is the visible part of the page. This has been used recently to investigate general reading patterns of online articles [6,30], including, for example, backtracking (scrolling back in the browser). The viewport is large, particularly when websites are accessed from desktop and laptop computers, so it can provide only a very general indication of the attention position.

The position of the mouse cursor can indicate the position of user attention more precisely [12]. When a user moves or clicks the mouse, the position of the mouse cursor is relatively close to the position of the eye gaze on the screen [4,28]. Mouse-tracking information, as an indication of user attention, has been studied in various contexts, including in web search [10,12,26,28], e-commerce [29], web marketing [32], performing tasks [24], and online surveys [3].

Collective user attention on a web page, based on mouse activity, can be visualized by heatmaps [17,20,21]. Such heatmaps are used in many commercial web analytics services [17].

Recent studies show that in the context of textual web pages, horizontal mouse movements in the direction of reading (left to right in English) are more frequent than mouse movements in other directions [16,19] and that those mouse movements are indeed related to reading activity [16].

Jarodzka and Brand-Gruwel [14] proposed to separate eye-tracking research on reading behaviors and patterns into three levels:

- Level 1: Reading sentences and single words;
- Level 2: Reading and comprehending a complete text;
- Level 3: Reading and processing several text documents.

Interestingly, while implicit client-side indicators are widely used as a replacement for eye-tracking in many applications, including in analysis of reading behaviors at levels 2 and 3 (a whole text and multiple documents, respectively), to the best of our knowledge, client-side indicators, including mouse movements, have never been used to analyze reading behaviors at level 1, i.e. sentences and words, which is what this paper does.

3 Mouse Cursor Tracking and Visualization

This study examined PAR activity of users on the ObjectDB developer guide pages.[1] A special PAR Server was implemented in order to track, store, report, and visualize mouse movements. Figure 1 shows the architecture that was used.

To track mouse movements, a reference to a *Tracking Script* was embedded in all the web pages. As a result, every page request from the website returned a revised version of the page that triggered an additional request to load the Tracking Script from the PAR Server. The Tracking Script (implemented in JavaScript) tracked the user's mouse move events and reported them back to the *Collector* component in the PAR Server. The Collector stored the data anonymized in a dedicated database, adhering to industry standards of data anonymization and user privacy preservation.

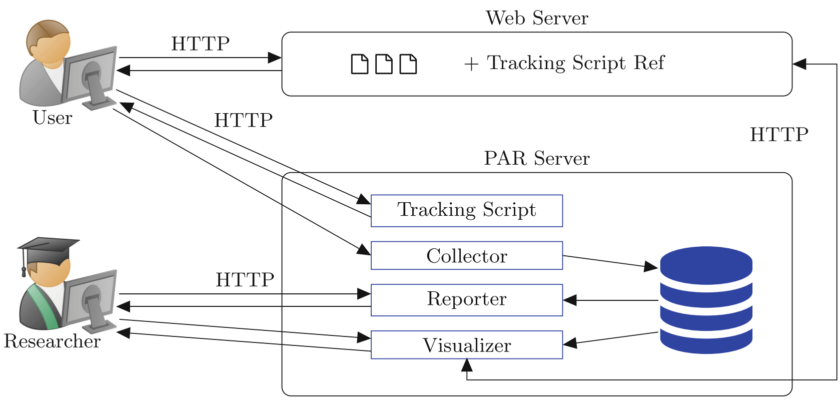

Fig. 1. Architecture of the PAR Server

Mouse movement data were available for research through two components in the PAR server. The *Reporter* component was used to retrieve mouse movement data, including cumulative statistics, via queries. The *Visualizer* component was used to visualize individual page views by displaying mouse movements on web pages. Figure 2 demonstrates the implemented visualization with a sample PAR activity of a real user.

Mouse movements to the right are displayed as solid green lines. Mouse movements to the left are displayed as dashed red lines. Different colors and styles for different directions are very helpful in inferring the mouse movement directions and the reading progression from a still picture. In addition, one direction represents mainly reading activity (the right direction for LTR languages such as English), and the other direction represents mainly progressing to the next line of text (the left direction for English), so distinct colors and styles help in identifying reading patterns.

[1] https://www.objectdb.com/java/jpa.

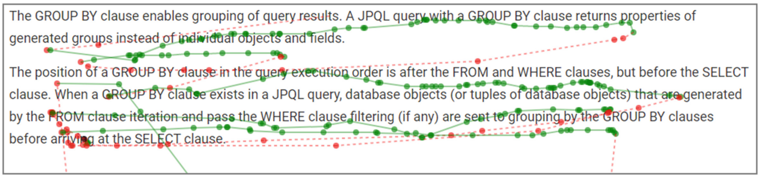

Fig. 2. Visualization of mouse movements for PAR activity analysis (Color figure online)

The circles on both the green and red lines represent the positions of the mouse cursor in reported mouse events. Mouse events are tracked at a rate of one event per one-tenth of a second (i.e. 10 events per second), so the green and red lines connecting adjacent circles represent mouse movements during one-tenth of a second.

4 Mining PAR Views

Unlike page scrolling, which is inevitable when reading web pages (and therefore, employed by all users), using the mouse cursor to point at sentences and words while reading is an uncommon reading pattern [16]. In some sense, it is a hidden phenomenon, and the likelihood of seeing PAR activity similar to that in Fig. 2, by exploring the mouse activity in a random page view, is low. This may explain why it had hardly been studied, despite the intensive work on using mouse movement data as an alternative to eye-tracking.

Note, that low prevalence does not mean that it is not useful, as feedback data from sample users is also beneficial for analytics purposes. However, low frequency means that to examine PAR, we need a method for mining potential PAR views in a web usage dataset, i.e. filtering page views with a high probability of PAR activity.

The dataset that was used in this study consists of 389,424 page views of the 69 web pages of the ObjectDB developer guide. Data were collected from desktop users for a period of several months (ending in March 2020). Mobile users have been excluded as they normally do not use a mouse. We refer to these page views as collection 1. Two filters have been applied on collection 1 in a row, generating collection 2 and collection 3, which are much smaller but have higher probabilities to include PAR activity.

The first filter is trivial. PAR activity requires moving the mouse. Therefore, PAR is more likely in page views where the users moved the mouse longer. Figure 3 shows the distribution of the page views in collection 1 by the total time of mouse movements in seconds. By focusing on page views that include at least 64 seconds of mouse movements, which we refer to as collection 2, we significantly increase the probability of observing PAR activity.

Naturally, we cannot cover every PAR session by ignoring 98.5% of collection 1, but our goal is modest, we prioritize precision over recall and can toler-

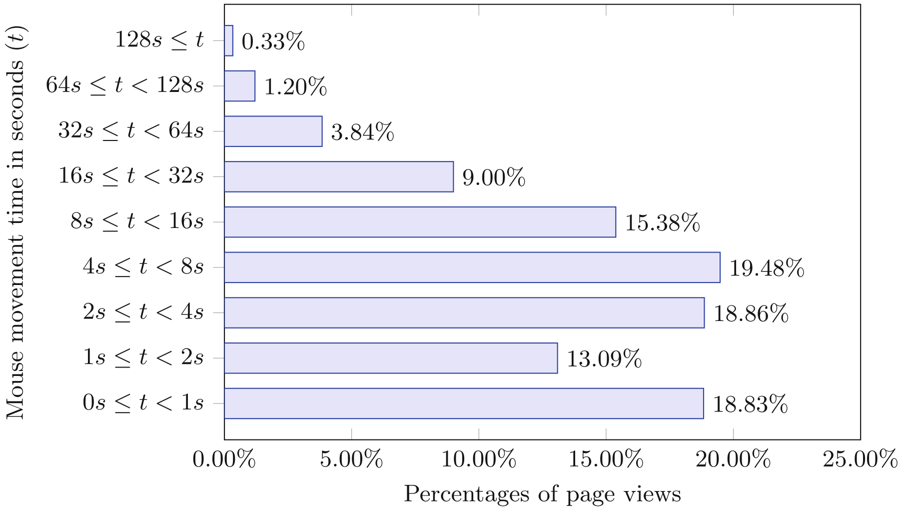

Fig. 3. Distribution of page views in collection 1 by mouse movement time

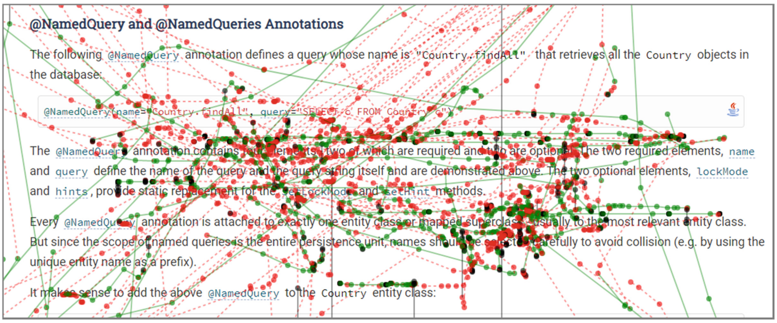

Fig. 4. Scattered mouse activity

ate a high false-negative rate as the objective is to find sample PAR sessions to investigate.

Collection 2 still has page views with no PAR activity, because frequent and long mouse movements are not necessarily related to PAR, as shown in Fig. 4.

On a side note, several similar examples of scattered mouse movements have been observed in collection 2. It is possible that the user fidgeted with the mouse while reading. It might also correlate with obsessive-compulsive disorders, but that is a matter for separate studies.

A second filter was used to retain only page views with considerably more mouse movement time to the right than mouse movement time to the left. The following ratio is calculated for each page view:

$$rlRatio(pageView) = \frac{rightTime(pageView)}{leftTime(pageView)} \tag{1}$$

where $rightTime(pageView)$ and $leftTime(pageView)$ are the total mouse movement times to the right and left for $pageView$, respectively.

The reasoning for using the second filter can be explained using Fig. 2. When the mouse is moved during PAR, it is moved in the direction of reading (to the right in this dataset, as the website is in English) at reading speed. It is moved much faster in the other direction, to the beginning of the next line of text. Therefore, a high $rlRatio(pageView)$ value in LTR texts (e.g. in English) may indicate that the mouse movements are not random but hint at PAR activity (and the opposite is true for RTL texts, such as in Hebrew and Arabic). These hypotheses regarding the directions and movements of the mouse during reading activity are supported by a recent statistical analysis of mouse movements [16]. Figure 5 shows the distribution of the page views in collection 2 by *rlRatio* values.

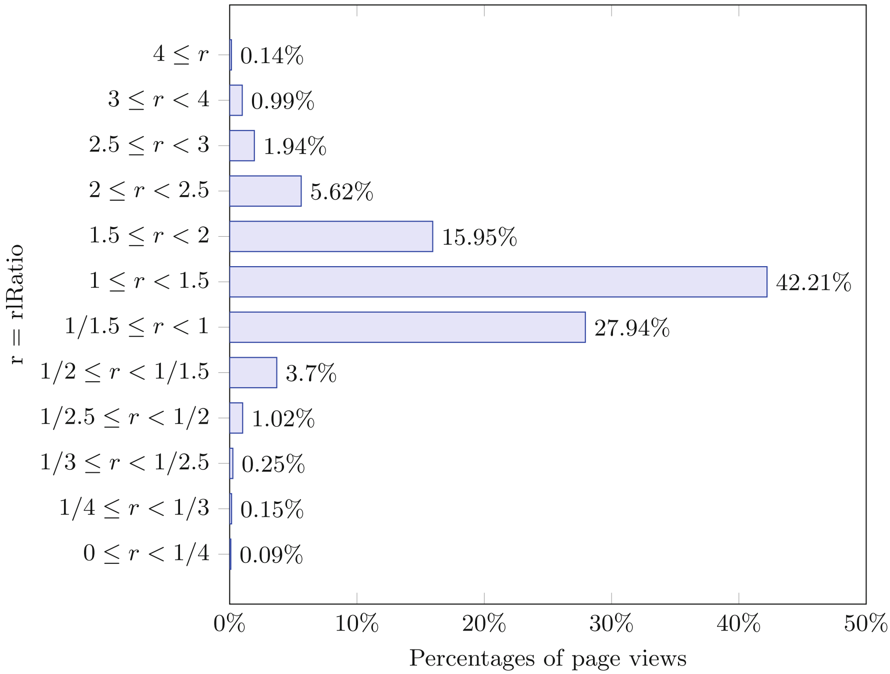

Fig. 5. Distribution of page views in Collection 2 by *rlRatio*

Note that the distribution of page views in Fig. 5 is not exactly symmetric around $rlRatio = 1$. As expected, $rlRatio \geq 1$ is significantly more frequent than $rlRatio < 1$ in this population of page views, which relative to collection 1, is richer with PAR activity.

To minimize false-positive errors (at the cost of a higher false-negative rate) and to improve the quality of the remaining page views (with respect to their PAR activity potential), this new filter was applied as a second-tier on top of the first filter, i.e. on collection 2. The result, collection 3, retains only page views in collection 2 with $rlRatio(pageView) \geq 3$, which is a small fraction of collection 1 (the complete dataset).

A computer program selected 40 random sample page views from each of the three collections and combined them into one set of 120 untagged sample page views. An examiner (an undergraduate student of International Business and Languages) was asked to count the number of text lines in each page view that appear to have been read with moving the mouse cursor as a pointer, i.e. hinting at PAR activity. Table 1 and Fig. 6 show the results. In Fig. 6, the 40 sample page views of each collection are in descending order of the number of identified PAR lines.

Table 1. Collections & Results

	Collection 1	Collection 2	Collection 3
Filter	None	$move\ time \geq 64$ s	$move\ time \geq 64$ s and $rlRatio \geq 3$
Collection size (views)	389,424	6,971	79
Sample size (views)	40	40	40
Views with PAR activity	8	36	40
% Views with PAR activity	20%	90%	100%
Identified PAR lines	16	217	481
Identified PAR lines/view	0.4	5.4	12.0

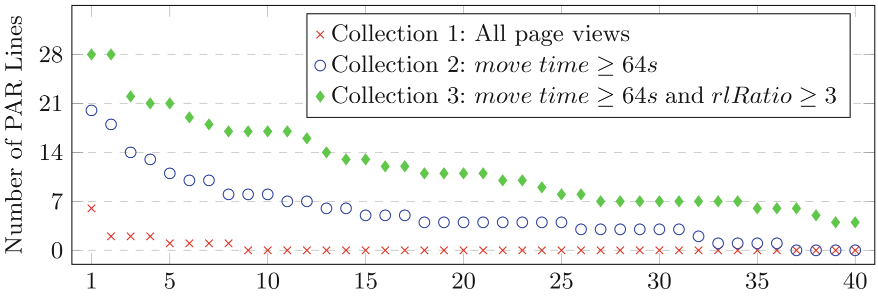

Fig. 6. Identified PAR lines in 120 page views

The sample page views from collection 1 had very little PAR activity (only 16 lines in total for 8 page views, where the other 32 page views had no identified PAR activity at all). This shows that PAR is indeed uncommon and that obtaining page views with PAR requires deliberate "mining". As expected, PAR activity was more frequent in collection 2, and even more so in collection 3.

5 Exploring PAR Using Examples

Mining page views with a high probability of PAR activity, as discussed in Sect. 4, combined with visualizing mouse cursor movements on pages, as described in Sect. 3, paves the way for exploring PAR by examining real examples.

Note that we can assume that the mouse movements visualized in these examples represent eye gaze movements (and accordingly, reading patterns) and that the positions of the mouse cursor reflect eye gaze positions. As discussed in Sect. 2, previous studies have already shown significant proximity between the mouse cursor position and the eye gaze position on a screen when a user moves the mouse.

The primary concern regarding the feasibility of using PAR data is that relevant mouse-tracking data may not be available, because as discussed previously, PAR is not a common reading technique. This section shows that the page views in collection 3 contain mouse-tracking data that are similar to eye-tracking data, and therefore, PAR-tracking data might replace eye-tracking data of web users, which are usually unavailable. Accordingly, the feasibility concern is reduced to whether the amount of PAR-tracking data that can be obtained is sufficient for practical applications. This is discussed in Sect. 6.

Figure 7 shows a typical PAR activity. Green, which represents movements from left to right, is more dominant than red, which represents movements in the opposite direction (from right to left). Lines between two adjacent circles represent mouse movements during one-tenth of a second. The green circles are much closer together than the red circles because right mouse movements are mostly carried out at reading speed, whereas left movements to the beginning of the next lines of text tend to be much faster [16].

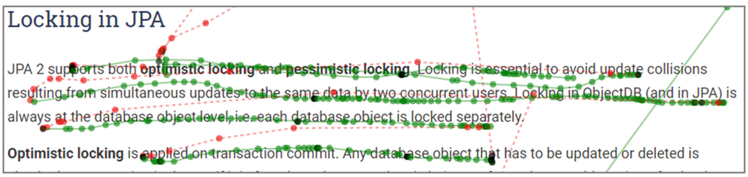

Fig. 7. A typical PAR activity (Color figure online)

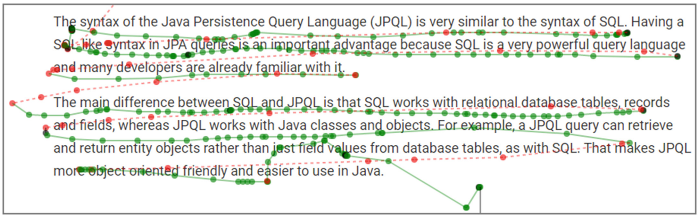

Fig. 8. Cutting corners at the end of text lines (Color figure online)

Most of the mouse cursor movements are shorter than the full-text width, because of a "cutting corners" effect, as shown, for example, in Fig. 7. Even users that use PAR extensively use shortcuts where possible, and apparently, moving the mouse to the very first and last characters of the text lines is not essential for benefiting from PAR. Some users cut corners on both the left and right sides of the text lines, as demonstrated in Fig. 7, while others move the mouse over the first word in text lines but occasionally skip the last word, as demonstrated in Fig. 8.

This is a drawback of PAR-tracking data compared to eye-tracking data, as eye-tracking provides accurate data also at the edges of text lines. To overcome this issue, the missing information regarding reading the start and the end of text lines has to be completed through other means. For example, further work could explore the possibility of extrapolating the missing information using the user's average reading speed.

The vertical position of the mouse cursor relative to the text lines varies. Sometimes it is above the text, sometimes below the text, and sometimes on the text. Changes in the vertical position of the mouse cursor during reading activity are probably due to the difficulty of tracing perfectly straight lines with manual movements.

In browsers, a mouse cursor placed on text usually shows as a thin vertical bar, and therefore, does not normally hide the characters in the text (unlike the green and red lines in the visualization that occasionally hide the text). Consequently, moving the cursor over the text, as seen in Fig. 7, does not necessarily disrupt reading.

The position of the mouse cursor relative to the text seems to reflect personal preferences. Figure 9 shows a PAR activity with the cursor position mainly below the text lines. At first glance, it may be challenging to figure out whether the user traces text lines from above or below. The division of the text into paragraphs is helpful in that regard, as usually the last line of a paragraph is shorter than the full-text width, so we can expect it to be matched with a relatively shorter green movement line.

Sometimes, however, as shown in Fig. 10, the mouse movement lines are vaguer, and their total count does not match the total number of text lines in the paragraph. In these cases, it is difficult to match mouse cursor movements

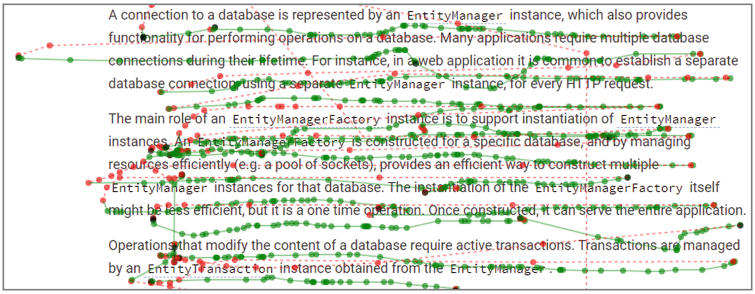

Fig. 9. Pointer below text: matching by text lines lengths

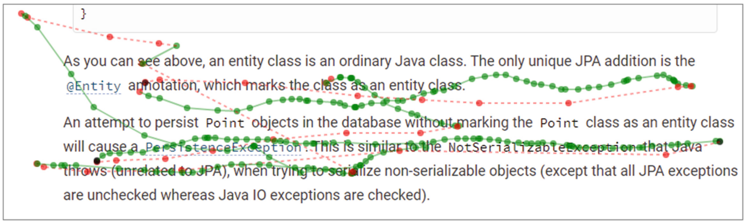

Fig. 10. Vague PAR: no clear way to match mouse cursor movements with text (Color figure online)

with text lines, and therefore, such mouse movement data may be less useful as PAR data.

When it is possible to match mouse movements with text words, it may provide valuable knowledge about the specific page view. When accumulated across distinct page views (of the same web page) of different users, it may also provide information about the page content.

An analysis of Fig. 11 gives several interesting insights. Firstly, we can find out which parts of the text have been read by the user and which parts have

Fig. 11. Skipping, repeating, and reading speed (Color figure online)

been skipped. In a page view with heavy PAR use, a jump such as the one from the word "lock" on the third line to the code fragment below, usually indicates that some lines of text have not been read. Although it is possible for a PAR user to occasionally read lines of text without moving the mouse, checking the timestamps of the mouse events before and after that jump (which are not shown in Fig. 11) confirms that, indeed, these text lines have been skipped by the user.

Secondly, we can see which parts of the text have been read more than once. Most of the first text line in Fig. 11 is covered by two distinct green lines, indicating that it has been reread.

Thirdly, the reading speed is revealed. The density of green circles on LockTimeoutException (on the first text line in Fig. 11), PESSIMISTIC_READ, and PESSIMISTIC_WRITE (on the last two text lines) is lower than the average density of green circles in this figure. Apparently, less time is needed to read these long technical strings (relative to their widths in pixels) compared to ordinary text words that require reading alongside thinking and comprehending. This demonstrates the dynamic nature of reading speed. Changes in reading speed are shown in most of the examples in this section. These might be valuable in spotting text complexity and unclarity.

Figure 12 shows an interesting attempt by a user to cut corners and read the last line quickly, but then, possibly because of difficulties in understanding the text at a glance, the user backtracks and reads that line again from the beginning.

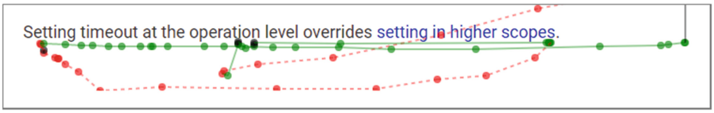

Fig. 12. Cutting corners and then rereading from the beginning (Color figure online)

Figure 13 shows another example of backtracking and rereading. A careful examination of this PAR activity (focusing on the green lines and circles) reveals that the first text line was read once, the second line, the third line, and the beginning of the fourth line were read twice, and the last sentence was read once. This visualization represents the reading activity of a single user. If this pattern repeats in the PAR activity of other users, it may indicate that these sentences are not sufficiently clear.

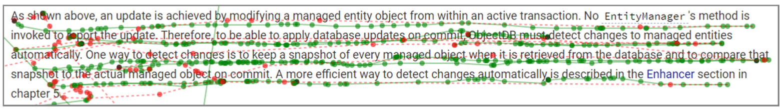

Fig. 13. Rereading lines of text (Color figure online)

Figure 14 shows another example of varying reading speed. The end of the sentence is read much slower than the beginning of the sentence. If this pattern repeats in the PAR activity of other users it may indicate that the end of the sentence is complex and may require simplification.

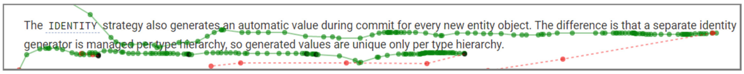

Fig. 14. Varying reading speed (Color figure online)

6 Discussion

The experiment discussed in Sect. 4 shows that PAR is not a common reading behavior. Based on the number of manually identified PAR lines in a sample of page views, as shown in Table 1, the frequency in the complete dataset (collection 1) is only 0.4 PAR lines per page view. This is a very rough estimate because of the small size of the sample (and we also expect to have different frequencies on different websites), but it shows that PAR activity is uncommon.

According to this estimate, there are over 155,000 PAR lines in this dataset, i.e. the examples in Sect. 5 cover only a tiny fraction of these PAR lines. Collecting eye-tracking on this scale could be complicated and expensive, so PAR-tracking data might be valuable. The 69 web pages used in this experiment, when presented in the most commonly used browser width on this website (1,920 pixels) contain 4,430 lines of text in total. Therefore, according to this estimate, each text line has on average about 35 PAR lines. The more frequently read lines (which are probably more important when considering how to improve the website content) may have much more PAR data in this dataset.

PAR-tracking might be a useful alternative to eye-tracking when eye-tracking data are unavailable. Particularly, eye-tracking may be impractical on public websites, as it requires active user collaboration, special equipment (to achieve high accuracy), and it may raise privacy concerns among users, as it makes use of cameras. On the other hand, PAR data can be collected in all modern browsers using embedded client-side JavaScript code.

Mouse movement tracking and recording is a core element of "session recording", which is a common practice in modern web analytics [17]. Session recording raises interesting questions in the context of user privacy and personal data protection, for example, due to the risk of collecting sensitive information through keystroke recording unintentionally [9]. However, session recording does not necessarily require prior user consent under personal data protection regulations, such as GDPR (under certain terms, as discussed by the IT and privacy lawyer

Arnoud Engelfriet [9]). If the data collected is completely anonymized, which is a standard practice in web analytics, then it is no longer considered personal data (e.g. by GDPR). Therefore, PAR-tracking can be used on almost any website, as long as data protection standards are preserved, and it could be useful if there are enough users and sufficient traffic.

This study has not used eye-tracking to confirm that the mouse movements that appear to be related to PAR activity indeed match the user's eye gaze while reading. As PAR activity is uncommon, such an experiment would require collecting eye-tracking data for a large number of users, in order to capture the natural PAR activity of the small percentage of users that use PAR. However, the significant proximity of the cursor position to the user's eye-gaze on a screen, when the user moves the mouse, has already been shown in previous studies, as discussed in Sect. 2. Based on this knowledge we can assume that mouse cursor movements along lines of text in the direction of reading (left to right in English) correlate with eye gaze movements in the same direction, and therefore, this reflects reading.

As shown in Sect. 5, the start and the end of text lines are often not covered by mouse movements, because of the “cutting corners” effect. This is a drawback of PAR data compared to eye-tracking data, and further work could look at practical solutions to complete the missing data by other means, for example, based on the user's average reading speed.

PAR-tracking data might be valuable in educational technology, e-learning, and web analytics, as it can show what users read, which parts of the text they skip, which sentences slow their reading speed down, which paragraphs they reread, etc., as demonstrated in Sect. 5. This information might help in spotting issues and obstacles in texts and in improving the overall quality of text content.

PAR-tracking may also be useful in research on reading, which is often performed using eye-tracking. Eliminating the availability and scalability constraints posed by eye-tracking and using PAR-tracking instead may open new doors. For example, international, worldwide research on reading behaviors and patterns on standard websites may become more feasible.

Because of the low prevalence of PAR, we cannot expect to analyze the behavior of every individual user. However, the applications mentioned above do not require data for every user. Data from a sample of users that use PAR might be sufficient in these applications, similar to the way that feedback from sample users is usually sufficient in surveys.

7 Conclusions and Further Work

This paper examines methods to extract and visualize PAR activity and demonstrates and analyzes sample PAR activity of real web users. The potential of PAR data in various applications, as well as the challenges, are discussed.

Further work could seek to develop PAR recognition methods to identify PAR more accurately, matching mouse movements to words in the text. PAR recognition ability might enable replacing eye-tracking with PAR-tracking in many applications, as discussed in this paper. PAR recognition could also help in exploring PAR further, e.g. in calculating the frequency of PAR more precisely and in learning about PAR users, which is essential if we want to treat PAR-tracking data as an accurate representation of the entire website audience.

References

1. Asvestopoulou, T., Manousaki, V., Psistakis, A., Andreadakis, V., Aslanides, I., Papadopouli, M.: DysLexML: screening tool for dyslexia using machine learning, pp. 1–6. ArXiv abs/1903.06274, March 2019
2. Biedert, R., Hees, J., Dengel, A., Buscher, G.: A robust realtime reading-skimming classifier. In: Proceedings of the Symposium on Eye Tracking Research and Applications, ETRA 2012, pp 123–130. Association for Computing Machinery, New York (2012). https://doi.org/10.1145/2168556.2168575
3. Cepeda, C., et al.: Mouse tracking measures and movement patterns with application for online surveys. In: Holzinger, A., Kieseberg, P., Tjoa, A.M., Weippl, E. (eds.) CD-MAKE 2018. LNCS, vol. 11015, pp. 28–42. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-99740-7_3
4. Chen, M.C., Anderson, J.R., Sohn, M.H.: What can a mouse cursor tell us more? correlation of eye/mouse movements on web browsing. In: CHI 2001 Extended Abstracts on Human Factors in Computing Systems, CHI EA 2001, pp. 281–282. Association for Computing Machinery, New York (2001)
5. Claypool, M., Le, P., Wased, M., Brown, D.: Implicit interest indicators. In: Proceedings of the 6th International Conference on Intelligent User Interfaces, IUI 2001, pp. 33–40. Association for Computing Machinery, New York (2001). https://doi.org/10.1145/359784.359836
6. Conlen, M., Kale, A., Heer, J.: Capture & analysis of active reading behaviors for interactive articles on the web. Comput. Graph. Forum **38**(3), 687–698 (2019). https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13720https://doi.org/10.1111/cgf.13720
7. Dirix, N., Vander Beken, H., De Bruyne, E., Brysbaert, M., Duyck, W.: Reading text when studying in a second language: an eye-tracking study. Read. Res. Q. (2019). https://doi.org/10.1002/rrq.277
8. Eraslan, S., Yesilada, Y., Harper, S.: "The best of both worlds!": integration of web page and eye tracking data driven approaches for automatic AOI detection. ACM Trans. Web **14**(1), 1–31 (2020)
9. Gilliam Haije, E.: Are session recording tools a risk to internet privacy (2018). https://mopinion.com/are-session-recording-tools-a-risk-to-internet-privacy/
10. Guo, Q., Agichtein, E.: Exploring mouse movements for inferring query intent. In: Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2008, pp. 707–708. Association for Computing Machinery, New York (2008). https://doi.org/10.1145/1390334.1390462

11. Hernandez, M., Wang, Y., Sheng, H., Kalliny, M., Minor, M.: Escaping the corner of death? an eye-tracking study of reading direction influence on attention and memory. J. Consum. Mark. **34**, 1–10 (2017). https://doi.org/10.1108/JCM-02-2016-1710
12. Huang, J., White, R., Buscher, G.: User see, user point: gaze and cursor alignment in web search. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2012, pp. 1341–1350. Association for Computing Machinery, New York (2012)
13. Iqbal, S., Bailey, B.: Using eye gaze patterns to identify user tasks. In: The Grace Hopper Celebration of Women in Computing, vol. 04, January 2004
14. Jarodzka, H., Brand-Gruwel, S.: Tracking the reading eye: towards a model of real-world reading. J. Comput. Assist. Learn. **33**(3), 193–201 (2017). https://onlinelibrary.wiley.com/doi/abs/10.1111/jcal.12189. https://doi.org/10.1111/jcal.12189
15. Javed, Y., Shehab, M.: Look before you authorize: Using eye-tracking to enforce user attention towards application permissions. Proc. Priv. Enhancing Technol. **2017**(2), 23–37 (2017). https://content.sciendo.com/view/journals/popets/2017/2/article-p23.xml
16. Kirsh, I.: Directions and speeds of mouse movements on a website and reading patterns: a web usage mining case study. In: Proceedings of the 10th International Conference on Web Intelligence, Mining and Semantics (WIMS 2020), Biarritz, France, pp. 129–138. Association for Computing Machinery, New York, June 2020. https://doi.org/10.1145/3405962.3405982
17. Kirsh, I., Joy, M.: A different web analytics perspective through copy to clipboard heatmaps. In: Bielikova, M., Mikkonen, T., Pautasso, C. (eds.) ICWE 2020. LNCS, vol. 12128, pp. 543–546. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-50578-3_41
18. Kirsh, I., Joy, M.: Splitting the web analytics atom: from page metrics and KPIs to sub-page metrics and KPIs. In: Proceedings of the 10th International Conference on Web Intelligence, Mining and Semantics (WIMS 2020), Biarritz, France, pp. 33–43. Association for Computing Machinery, New York, June 2020. https://doi.org/10.1145/3405962.3405984
19. Kirsh, I., Joy, M., Kirsh, Y.: Horizontal mouse movements (HMMs) on web pages as indicators of user interest. In: Proceedings of the 22nd HCI International Conference (HCII 2020), Communications in Computer and Information Science. Springer, Cham, July 2020
20. Lamberti, F., Paravati, G., Gatteschi, V., Cannavò, A.: Supporting web analytics by aggregating user interaction data from heterogeneous devices using viewport-DOM-based heat maps. IEEE Trans. Industr. Inf. **13**, 1989–1999 (2017)
21. Lamberti, F., Paravati, G.: VDHM: viewport-DOM based heat maps as a tool for visually aggregating web users' interaction data from mobile and heterogeneous devices. In: Proceedings of the 2015 IEEE International Conference on Mobile Services, MS 2015, USA, pp. 33–40. IEEE Computer Society (2015)
22. Lee, J., Ahn, J.H.: Attention to banner ads and their effectiveness: an eye-tracking approach. Int. J. Electron. Commer. **17**(1), 119–137 (2012). https://doi.org/10.2753/JEC1086-4415170105

23. Li, Y., Xu, P., Lagun, D., Navalpakkam, V.: Towards measuring and inferring user interest from gaze. In: Proceedings of the 26th International Conference on World Wide Web Companion, WWW 2017 Companion, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, pp. 525–533 (2017). https://doi.org/10.1145/3041021.3054182
24. Milisavljevic, A., Hamard, K., Petermann, C., Gosselin, B., Doré-Mazars, K., Mancas, M.: Eye and mouse coordination during task: from behaviour to prediction. In: International Conference on Human Computer Interaction Theory and Applications, Setúbal, Portugal, pp. 86–93. SciTePress, January 2018. https://doi.org/10.5220/0006618800860093
25. Müller, A.M., Baumgartner, S.E., Kühne, R., Peter, J.: The effects of social information on the enjoyment of online videos: an eye tracking study on the role of attention. Media Psychol., 1–22 (2019). https://doi.org/10.1080/15213269.2019.1679647
26. Navalpakkam, V., Jentzsch, L., Sayres, R., Ravi, S., Ahmed, A., Smola, A.: Measurement and modeling of eye-mouse behavior in the presence of nonlinear page layouts. In: Proceedings of the 22nd International Conference on World Wide Web, WWW 2013, New York, NY, US, pp. 953–964. Association for Computing Machinery (2013). https://doi.org/10.1145/2488388.2488471
27. Pellicer-Sánchez, A.: Incidental vocabulary acquisition from and while reading: an eye-tracking study. Stud. Second Lang. Acquisition **36**(1), 97–130 (2015)
28. Rodden, K., Fu, X.: Exploring how mouse movements relate to eye movements on web search results pages. In: Proceedings of ACM SIGIR 2007 Workshop on Web Information Seeking and Interaction, pp. 29–32. Association for Computing Machinery, New York (2007). http://research.microsoft.com/~ryenw/proceedings/WISI2007.pdf
29. Schneider, J., Weinmann, M., vom Brocke, J., Schneider, C.: Identifying preferences through mouse cursor movements - preliminary evidence. In: Proceedings of the 25th European Conference on Information Systems (ECIS), Guimarães, Portugal, pp. 2546–2556. Research-in-Progress Papers (2017)
30. Smadja, U., Grusky, M., Artzi, Y., Naaman, M.: Understanding reader backtracking behavior in online news articles. In: The World Wide Web Conference, pp. 3237–3243. WWW 2019. Association for Computing Machinery, New York (2019). https://doi.org/10.1145/3308558.3313571
31. Strindberg, A.: Eye movements during reading and reading assessment in Swedish school children: a new window on reading difficulties. In: Proceedings of the 11th ACM Symposium on Eye Tracking Research & Applications, ETRA 2019. Association for Computing Machinery, New York (2019). https://doi.org/10.1145/3314111.3322878
32. Tzafilkou, K., Protogeros, N., Yakinthos, C.: Mouse tracking for web marketing: enhancing user experience in web application software by measuring self-efficacy and hesitation levels. Int. J. Strateg. Innov. Mark. **01** (2014). https://doi.org/10.15556/IJSIM.01.04.005

33. Zahoor, S., Bedekar, M., Kosamkar, P.K.: User implicit interest indicators learned from the browser on the client side. In: Proceedings of the 2014 International Conference on Information and Communication Technology for Competitive Strategies, ICTCS 2014. Association for Computing Machinery, New York (2014)
34. Zhang, H., Miller, K., Cleveland, R., Cortina, K.: How listening to music affects reading: evidence from eye tracking. J. Exp. Psychol. Learn. Mem. Cogn. 44 (2018). https://doi.org/10.1037/xlm0000544