Using Mouse Movement Heatmaps to Visualize User Attention to Words

Ilan Kirsh
kirsh@mta.ac.il
The Academic College of Tel Aviv-Yaffo
Tel Aviv, Israel

ABSTRACT

User attention on web pages is commonly visualized using heatmaps. Substantial mouse activity in an area of a web page may indicate substantial user attention. Accordingly, areas of web pages are painted with hot and cold colors based on the frequency of mouse clicks and movements in these areas. This paper introduces a new type of web page heatmap, the Word Attention Heatmap (WAH), which visualizes user attention to text words based on mouse movements. Naturally, complex words and unclear text may draw more attention and increase reading time, and therefore, may be painted in hotter colors. Consequently, WAHs may help in identifying complex words and challenging sentences as part of a process of improving and simplifying textual web content.

CCS CONCEPTS

• **Human-centered computing** → **Heat maps**; **Pointing devices**; • **Information systems** → **Browsers**.

KEYWORDS

Visualization, Heatmap, Website, Web Analytics, Mouse Cursor, Pointer Assisted Reading (PAR), User Behavior, Text Simplification, Complex Word Identification (CWI).

ACM Reference Format:
Ilan Kirsh. 2020. Using Mouse Movement Heatmaps to Visualize User Attention to Words. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society (NordiCHI '20), October 25–29, 2020, Tallinn, Estonia.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3419249.3421250

1 INTRODUCTION

Web analytics is essential for understanding the interests, preferences, needs, and actions of website users. The effectiveness of web analytics has been demonstrated in a wide range of industries, including online learning [21], online news [11], e-commerce [9], and digital marketing [3, 12]. Web analytics concepts, principles, and methods are described in detail in various books [1, 5, 13, 14].

During mouse activity (movement and clicks) there is a correlation between the position of the mouse cursor on the screen and the user's eye gaze [4, 10, 25]. Consequently, the mouse cursor position is often used to estimate which areas of a web page capture the user's attention. This has been used successfully in various applications, including online surveys [2], task execution [23], e-commerce [26], web marketing [28], and web search [8, 10, 24, 25].

The cumulative attention of all the visitors can be visualized using heatmaps [19, 20, 29]. Areas with high mouse activity are displayed with hot background colors (e.g. red) and areas with low mouse activity with cold background colors (e.g. blue). Different shades of hot and cold colors represent different levels of activity. The visualization can help website maintainers in improving and optimizing the structure and the user interface of websites. Attention heatmaps, based on mouse activity, are offered by various commercial web analytics services [17].

This paper introduces a new type of heatmap, the Word Attention Heatmap (WAH), with a different goal. Instead of visualizing user attention to web page areas, WAHs visualize the user's attention to specific text words while reading, highlighting words that require longer reading time. Unlike a hot region in a standard web page heatmap, which is usually (not always) a positive sign of user interest, a "hot" word on a WAH may indicate user difficulties in reading (similarly to frequent copying of a word to the clipboard, which may also indicate complexity [15]). Identifying complex words that challenge users can help in improving content and is the first step in text simplification processes [27].

2 IMPLEMENTATION ARCHITECTURE

Figure 1 shows the architecture of a WAH implementation. The client-side is on the left, and the server-side is on the right. At the top, we can see a standard HTTP client-server communication between a browser and a web server.

To track mouse activity, a reference to a *Tracking Script* is embedded in all the relevant website pages. As a result, loading a web page from the web server is followed by loading the *Tracking Script* from the WAH server. The script tracks mouse events and reports them to the *Collector* component in the WAH server, which stores the data (following anonymization) in a dedicated database. Mouse event data include event types (mouse move, click, etc.), event times, and event positions (x,y) in pixels on the web pages. Mouse move events are sampled at a rate of up to 10 events per second. In addition, the script maps text words to positions in pixels, and that mapping is also reported and stored in the database. This is required in order to link mouse events to words because the layout of words on a web page is client dependent (e.g. if the browser window is wider, text lines can be longer and contain more words).

To see a WAH, a web analyst visits a web page through the *Visualizer* component, which serves as a proxy server. The Visualizer retrieves original web pages from the web server, paints them with heatmap colors based on the mouse activity data in the database

NordiCHI '20, October 25–29, 2020, Tallinn, Estonia

ACM ISBN 978-1-4503-7579-5/20/10.
https://doi.org/10.1145/3419249.3421250

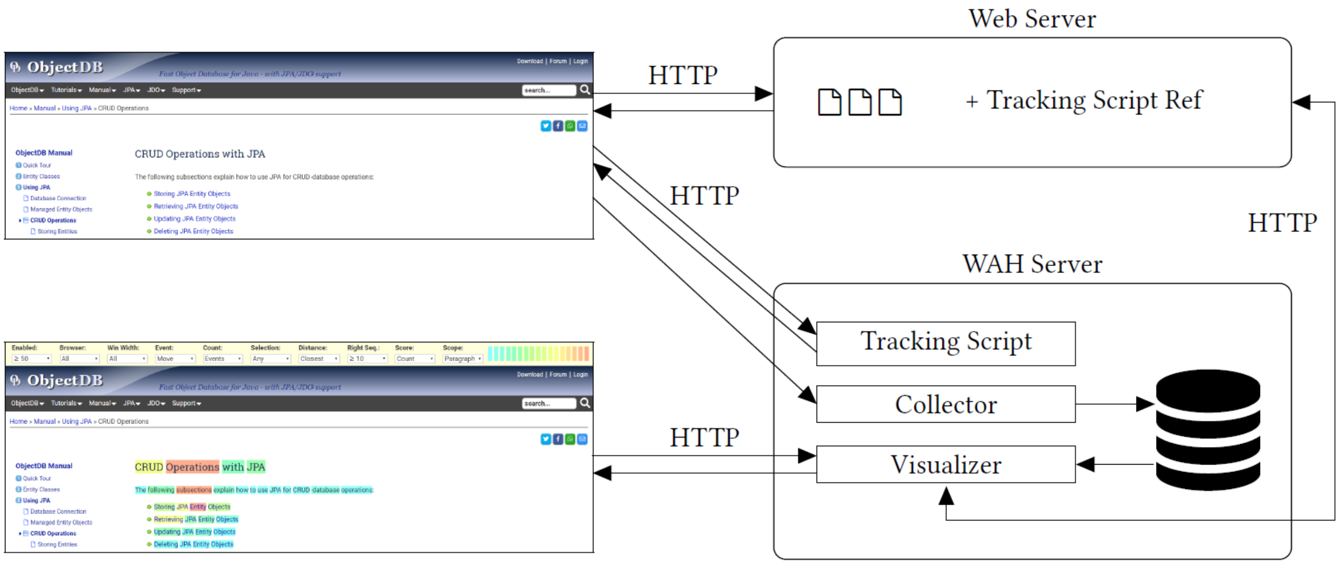

Figure 1: High-Level Architecture of the Word Attention Heatmap (WAH) Implementation

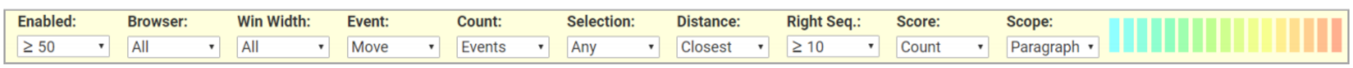

Figure 2: Word Attention Heatmap (WAH) Control Toolbar

(and the current settings), and then returns the modified heatmap pages to the analyst's browser with an additional settings toolbar at the top. Any setting change is reported back to the Visualizer and an updated heatmap, based on the new settings, is created by the Visualizer and loaded by the browser.

3 WORD SCORING AND SETTINGS

This section explains how the Visualizer component (see Figure 1) ranks words and sets their colors. We start with a description that is based on the default settings (which are used in all the examples in section 4 below).

The score of a word is the number of mouse move events that are associated with it. Only mouse move events that represent movements to the right and are part of sequences of at least 10 mouse move events to the right are counted. The motivation is to focus on mouse movements that may be related to Pointer Assisted Reading, which is reading text with the aid of the mouse (as defined in [16, 18]), i.e. from left to right in English (the website language). Every mouse event in the text area is associated with the closest text word. The heatmap is disabled for paragraphs in which the maximum word score is less than 50 (insufficient amount of data). Figure 2 shows the settings toolbar. The legend on the right side presents the fixed palette of colors from cold (left) to hot (right). The word with the highest score in the paragraph receives the hottest color, whose index is 16 (red), and that sets the scale for the other words in that paragraph, which receive colors proportionally to their scores. For example, a word with half of the highest score receives the color with index 8 (green). Unlike ordinary web page attention heatmaps, in WAH every paragraph is processed separately from other paragraphs, to neutralize differences in popularity of paragraphs.

Many of the default settings that are described above are configurable using the settings toolbar, including the threshold score that enables heatmaps (50 by default) and the right sequence threshold (minimum 10 events by default). The toolbar supports filtering data by browser type (e.g. Chrome) and client window width (e.g. 1920 pixels). Other options include visualizing mouse clicks and text selections, counting multiple events on the same page view only once, different methods for associating mouse events with words (e.g. associating a mouse event with multiple words within a radius), and various score normalization methods. The results are relatively stable with most of the settings. The default settings are somewhat arbitrary, and further work is needed to explore the effect of different options. There is also an option for ranking words relative to the page rather than to the paragraph, generating heatmaps that are more similar to ordinary web page attention heatmaps.

4 DEMONSTRATION

WAHs are demonstrated in this section on pages of the ObjectDB Developer Guide (www.objectdb.com/java/jpa). Mouse tracking data were collected for a period of three months ending in March 2020 for all visitors.

The examples in this section are from pages that had between 6,049 to 26,674 views. The score of the hottest words (i.e. number

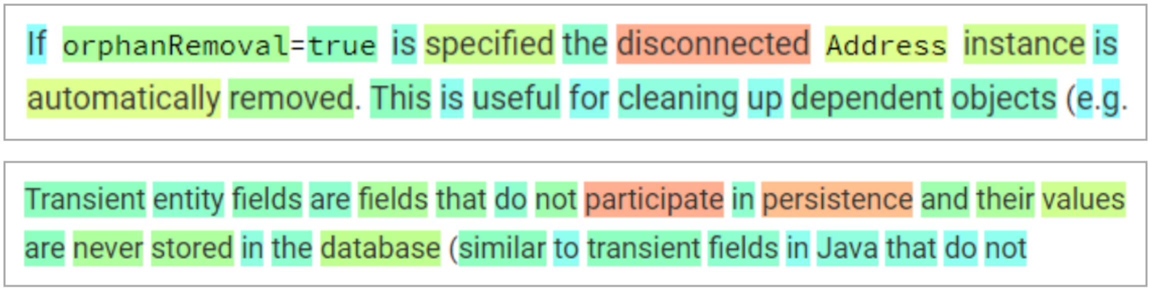

Figure 3: Complex Words: "disconnected" (2234), "participate" (3189), and "persistence" (3158)

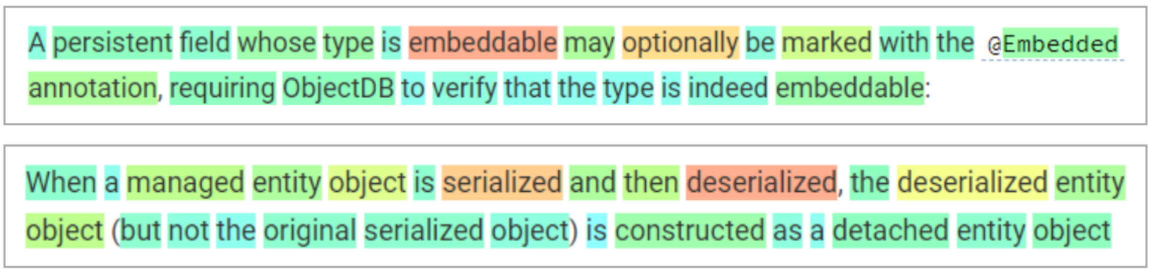

Figure 4: Complex Technical Words: "embeddable" (563), "serialized" (1085), and "deserialized" (1223)

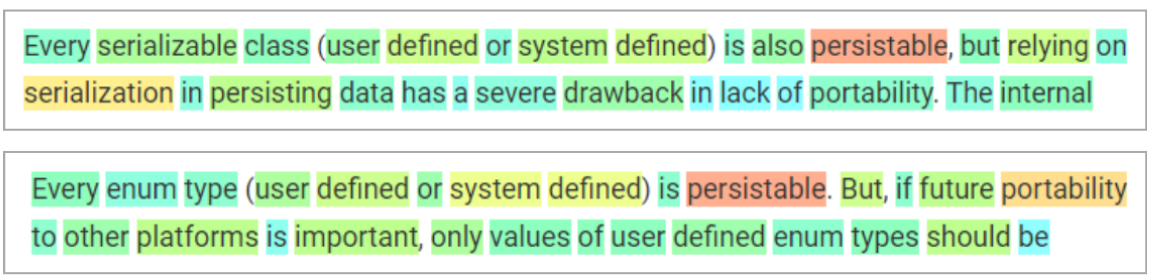

Figure 5: Another Complex Word: "persistable" (232, 96)

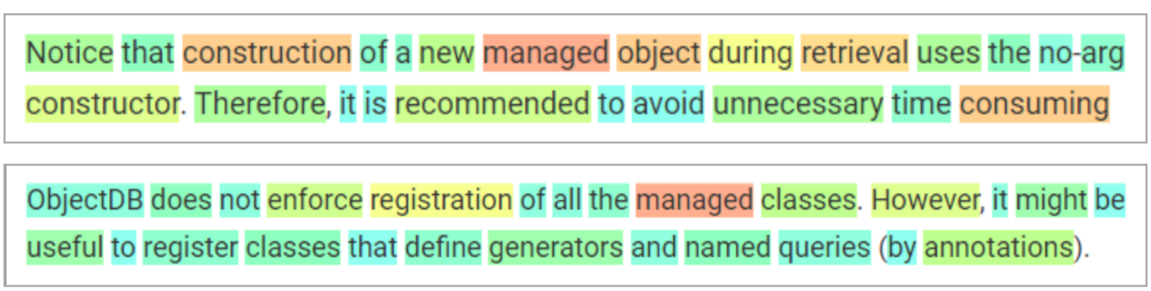

Figure 6: A Simple Word with a Challenging Meaning in This Context: "managed" (374, 132)

of mouse move events) is shown in parentheses in the captions of Figures 3, 4, 5, and 6 and is considerably lower than the number of views. This is expected, as Pointer Assisted Reading (reading with the aid of the mouse) is occasional [16, 18]. Moreover, most visitors do not read whole pages, so not all page views feed into the scores of each individual word in the page.

Words with more syllables are considered complex in classic readability tests such as the Gunning fog index [7] and the SMOG grade [22] (this rule has exceptions, e.g. "Wikipedia" is not a complex word for most web users). It is easy to see in the following examples that short words are indeed "colder". Each one of Figures 3, 4, 5, and 6 shows two independent heatmaps of separate paragraphs. The focus in each heatmap is on the "hottest" word in the paragraph (colored red) and its close environment. Hot words, with substantial user attention relative to other words in their close environment (based on mouse move events), demonstrate slower average reading speed and potentially higher complexity to users.

In Figure 3, the words "disconnected", "participate", and "persistence" are marked as hot. These are long words with 3-4 syllables. Note that the longest word, "automatically" (6 syllables), attracts less user attention, as well as the words "dependent", "database", "entity", and "specified" (3 syllables), which are frequently used in this developer guide. These words may be less challenging to users of this website.

The technical terms, "embeddable", "serialized", and "deserialized" in Figure 4 are colored red. These terms are explained on the website but, apparently, some users find them challenging (for example, users that arrive at these paragraphs directly). Note that each one of these words appears twice in these examples. The first occurrence is hot (red), and the second is colder (yellow / green). Reading the same word again in the same paragraph is expected to be easier.

The word "persistable" is another complex technical word. It appears several times in the examined web pages, and most of its occurrences (Figure 5 shows two of them) are colored red.

The word "managed" (2 syllables, #2,591 in the list of the most frequent words in Google's Trillion Word Corpus [6]), is a relatively simple word. However, in this context (see Figure 6), it may be unclear. Readers may ask themselves what is a "managed object" and what are "managed classes". These are, again, technical terms that are explained elsewhere in the developer guide. The last examples are particularly important as they show that tracking the behavior and the mouse activity of real users can reveal text complexity, which could be more difficult to find using static text analysis.

5 CONCLUSIONS AND FURTHER WORK

The new proposed heatmaps highlight words that attract more mouse movements than other words in their paragraphs. Previous works show that cumulative mouse movements reflect cumulative user attention and that some users read with the aid of the mouse. Excess attention and reading time may indicate potential issues in the readability and understandability of the text. The examples in this paper support this idea visually. Further work should build on the visual aspects that this paper presents by carrying out quantitative research of this user behavior. The ability to identify complex text in websites could be useful in the process of content improvement and text simplification.

REFERENCES

[1] Feras Alhlou, Shiraz Asif, and Eric Fettman. 2016. *Google Analytics Breakthrough: From Zero to Business Impact.* John Wiley & Sons, USA.

[2] Catia Cepeda, Joao Rodrigues, Maria Camila Dias, Diogo Oliveira, Dina Rindlisbacher, Marcus Cheetham, and Hugo Gamboa. 2018. Mouse Tracking Measures and Movement Patterns with Application for Online Surveys. In *Machine Learning and Knowledge Extraction*, Andreas Holzinger, Peter Kieseberg, A Min Tjoa, and Edgar Weippl (Eds.). Springer International Publishing, Cham, 28–42.

[3] Dave Chaffey and Mark Patron. 2012. From web analytics to digital marketing optimization: Increasing the commercial value of digital analytics. *Journal of Direct, Data and Digital Marketing Practice* 14 (07 2012), 30–45. https://doi.org/10.1057/dddmp.2012.20

[4] Mon Chu Chen, John R. Anderson, and Myeong Ho Sohn. 2001. What Can a Mouse Cursor Tell Us More? Correlation of Eye/Mouse Movements on Web Browsing. In *CHI '01 Extended Abstracts on Human Factors in Computing Systems* (Seattle, Washington) *(CHI EA '01)*. Association for Computing Machinery, New York, NY, USA, 281–282. https://doi.org/10.1145/634067.634234

[5] Brent Dykes. 2014. *Web Analytics Kick Start Guide: A Primer on the Fundamentals of Digital Analytics.* Peachpit, Pearson Education, USA.

[6] first20hours. 2020. 20k.txt, most common English words in order of frequency, as determined by n-gram frequency analysis of the Google's Trillion Word Corpus. https://github.com/first20hours/google-10000-english

[7] Robert Gunning. 1952. *The Technique of Clear Writing.* McGraw-Hill, New York, New York.

[8] Qi Guo and Eugene Agichtein. 2008. Exploring Mouse Movements for Inferring Query Intent. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Singapore, Singapore) *(SIGIR '08)*. Association for Computing Machinery, New York, NY, USA, 707–708. https://doi.org/10.1145/1390334.1390462

[9] Layla Hasan, Anne Morris, and Steve Probets. 2009. Using Google Analytics to Evaluate the Usability of E-Commerce Sites. In *Proceedings of the 1st International Conference on Human Centered Design.* Springer Berlin Heidelberg, Berlin, Heidelberg, 697–706. https://doi.org/10.1007/978-3-642-02806-9_81

[10] Jeff Huang, Ryen White, and Georg Buscher. 2012. User See, User Point: Gaze and Cursor Alignment in Web Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. Association for Computing Machinery, New York, NY, USA, 1341–1350. https://doi.org/10.1145/2207676.2208591

[11] Edson C. Tandoc Jr. 2015. Why Web Analytics Click. *Journalism Studies* 16, 6 (2015), 782–799. https://doi.org/10.1080/1461670X.2014.946309 arXiv:https://doi.org/10.1080/1461670X.2014.946309

[12] Joel Järvinen and Heikki Karjaluoto. 2015. The use of Web analytics for digital marketing performance measurement. *Industrial Marketing Management* 50 (2015), 117 – 127. https://doi.org/10.1016/j.indmarman.2015.04.009

[13] Avinash Kaushik. 2007. *Web Analytics: An Hour a Day.* SYBEX Inc., USA.

[14] Avinash Kaushik. 2010. *Web Analytics 2.0.* SYBEX Inc., USA.

[15] Ilan Kirsh. 2020. Automatic Complex Word Identification Using Implicit Feedback From User Copy Operations. In *Proceedings of the 21st International Conference on Web Information Systems Engineering (WISE 2020), Lecture Notes in Computer Science, forthcoming.* Springer International Publishing, Cham, 12.

[16] Ilan Kirsh. 2020. 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.* Association for Computing Machinery, New York, NY, USA, 129–138. https://doi.org/10.1145/3405962.3405982

[17] Ilan Kirsh and Mike Joy. 2020. A Different Web Analytics Perspective Through Copy to Clipboard Heatmaps. In *Proceedings of the 20th International Conference on Web Engineering (ICWE 2020), Lecture Notes in Computer Science, vol 12128.* Springer International Publishing, Cham, 543–546. https://doi.org/10.1007/978-3-030-50578-3_41

[18] Ilan Kirsh and Mike Joy. 2020. Exploring Pointer Assisted Reading (PAR): Using Mouse Movements to Analyze Web Users' Reading Behaviors and Patterns. In *Proceedings of the 22nd HCI International Conference (HCII 2020), Lecture Notes in Computer Science.* Springer International Publishing, Cham, 18.

[19] Fabrizio Lamberti and Gianluca Paravati. 2015. 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 '15).* IEEE Computer Society, USA, 33–40. https://doi.org/10.1109/MobServ.2015.15

[20] F. Lamberti, Gianluca Paravati, Valentina Gatteschi, and Alberto Cannavò. 2017. Supporting Web Analytics by Aggregating User Interaction Data From Heterogeneous Devices Using Viewport-DOM-Based Heat Maps. *IEEE Transactions on Industrial Informatics* 13 (08 2017), 1989 – 1999. Issue 4. https://doi.org/10.1109/TII.2017.2658663

[21] Heng Luo, Stevie Rocco, and Carl Schaad. 2015. Using Google Analytics to Understand Online Learning: A Case Study of a Graduate-Level Online Course. In *Proceedings of the 2015 International Conference of Educational Innovation*

through Technology (EITT '15). IEEE Computer Society, USA, 264–268. https://doi.org/10.1109/EITT.2015.62

[22] G. H McLaughlin. 1969. SMOG grading: A new readability formula. *Journal of Reading* 12 (08 1969), 639–646. Issue 8.

[23] Alexandre Milisavljevic, Kevin Hamard, Coralie Petermann, Bernard Gosselin, Karine Doré-Mazars, and Matei Mancas. 2018. Eye and Mouse Coordination During Task: From Behaviour to Prediction. In *International Conference on Human Computer Interaction Theory and Applications* (Funchal, Madeira, Portugal). SciTePress, Setúbal, Portugal, 86–93. https://doi.org/10.5220/0006618800860093

[24] Vidhya Navalpakkam, LaDawn Jentzsch, Rory Sayres, Sujith Ravi, Amr Ahmed, and Alex Smola. 2013. 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* (Rio de Janeiro, Brazil) *(WWW '13)*. Association for Computing Machinery, New York, NY, USA, 953–964. https://doi.org/10.1145/2488388.2488471

[25] Kerry Rodden and Xin Fu. 2007. 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*. Association for Computing Machinery, New York, NY, USA, 29–32. http://research.microsoft.com/~ryenw/proceedings/WISI2007.pdf

[26] Johannes Schneider, Markus Weinmann, Jan vom Brocke, and Christoph Schneider. 2017. Identifying Preferences Through Mouse Cursor Movements – Preliminary Evidence. In *Proceedings of the 25th European Conference on Information Systems (ECIS)*. Research-in-Progress Papers, Guimarães, Portugal, 2546–2556.

[27] Matthew Shardlow. 2014. A Survey of Automated Text Simplification. *International Journal of Advanced Computer Science and Applications(IJACSA), Special Issue on Natural Language Processing 2014* 4, 1 (2014), 58–70. https://doi.org/10.14569/SpecialIssue.2014.040109

[28] Katerina Tzafilkou, Nicolaos Protogeros, and C.G. Yakinthos. 2014. Mouse Tracking for Web Marketing: Enhancing User Experience in Web Application Software by Measuring Self-Efficacy and Hesitation Levels. *International Journal on Strategic Innovative Marketing* 01 (2014). https://doi.org/10.15556/IJSIM.01.04.005

[29] Oleg Špakov and Darius Miniotas. 2007. Visualization of eye gaze data using heat maps. *Elektronika ir Elektrotechnika - Medicine Technology* 115 (01 2007).