Analysis of User Behavior in Interfaces with Recommended Items: An Eye-tracking Study

Peter Gaspar
Slovak University of Technology in Bratislava,
Faculty of Informatics and Information Technologies
Bratislava, Slovakia
name_surname@stuba.sk

Michal Kompan
Slovak University of Technology in Bratislava,
Faculty of Informatics and Information Technologies
Bratislava, Slovakia
name.surname@stuba.sk

Jakub Simko
Slovak University of Technology in Bratislava,
Faculty of Informatics and Information Technologies
Bratislava, Slovakia
name.surname@stuba.sk

Maria Bielikova
Slovak University of Technology in Bratislava,
Faculty of Informatics and Information Technologies
Bratislava, Slovakia
name.surname@stuba.sk

ABSTRACT

When analyzing user implicit feedback in recommender systems, several biases need to be taken into account. A user is influenced by the position (i.e., position bias) or by the appeal of the items (i.e., visual bias). Since images have become an essential part of the Web, the study of their impact on user behavior during the decision-making tasks is fundamental. This work contributes to the understanding of attention bias in item lists interfaces of recommenders. We present an eye-tracking user study that strives to analyze users' behavior in the task of choosing a movie to watch. Items are shown to users using two alternative representations: textual and image. We found changes in the user's behavior when the image type of interface is present. Based on our findings, the visual appeal of the images made users to change their gaze sequences more frequently.

KEYWORDS

user feedback, visual bias, eye-tracking, recommendation

ACM Reference Format:
Peter Gaspar, Michal Kompan, Jakub Simko, and Maria Bielikova. 2018. Analysis of User Behavior in Interfaces with Recommended Items: An Eye-tracking Study. In *Proceedings of Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, October 2018, Vancouver, Canada (IntRS Workshop).* 5 pages.

1 INTRODUCTION

Interpretation of the feedback in recommender systems has been an open problem for many years. Many approaches in this field still suffer from the lack of satisfactory feedback interpretation on recommended items. While gathering an implicit feedback is easy (e.g., clicks on items), distinction between the positive and negative one is tough. By clicking on an item, the user implicitly expresses a positive feedback, though his/her attitude may change after learning more details about the item. On the other hand, *not clicking on an item* does not automatically mean that the user is not interested in it: he/she might not have seen it. To solve this ambiguity, we need the information on user's visual attention. Unfortunately, a visual attention can be reliably measured only by gaze-tracking, which is unavailable in practical scenarios. Therefore, instead of measuring visual attention, researchers try to model the typical visual attention patterns and predict the gaze behavior according to them [26].

The order of the gaze visits is influenced by *attention bias* – a tendency to look on certain item(s) earlier than on others [13, 18]. This may depend on many factors, such as user interface layout, visual style, or item content representation. For example, one-dimension item lists may induce different behavior than two-dimension ones. Or, the user can scan through textual items differently than through those represented by images or animations. Moreover, attention bias may also depend on user characteristics such as goals, skills or cultural background. There are two main types of attention biases: position bias (induced by the position of the items) and visual bias (induced by the visual appeal of the items).

In this work we studied the attention bias in recommended item lists, where we compared the user behavior in textual and image representation of the items. We proposed and conducted an eye-tracking user study, where participants had to choose a movie to watch. Movies were presented either with their title or poster in the circular layout (due to the eye-tracking methodological reasons, as explained later). Our main goal was to investigate the possible differences in gaze paths between items represented either by text, or image.

Our findings show that textual and graphical representations of items induce different participants behavior in their gaze sequences. Users tend to skip more items when viewing interface with the images and make bigger transitions between the items as well. Moreover, when viewing content based on the preferred genres, users tend to make smaller transitions than in case of random content. Results of our study support our assumptions that the items represented by images may change the users' attention and there is clearly a need to take them into account in better understanding of user behavior in recommender systems.

2 USER FEEDBACK INTERPRETATION

In recommender systems a user feedback is a fundamental part of the process of user modelling. It is used to better understand user's preferences about the items and this information is further used to train recommendation models. However, the user feedback can be also misinterpreted due to the various influences. Thus, it is beneficial to identify these influences and take them into account while analyzing and interpreting users' behavior.

There are several factors that may influence user's behavior while browsing the items on the Web, such as personal characteristics (most notably demography, personality, emotions, and mood [23]) or user's short-term goal [14]. However, one of the most important aspects that may influence the user are the *items* themselves and the way they are presented.

Recommended items are usually presented in a form of a list or a grid [3]. The research in the recommendation lists has been already well mapped among researchers [8, 10] with grids gaining a recent interest [26]. Here, the researchers analyze users' behavior in lists/grids and utilize it to predict, which item would user prefer [2, 6, 26]. These methods take into account users' clicks in the ranked lists/grids and assume that the clicked item is also relevant for a user.

Chen et al. compared user behavior in three recommender interfaces [5]: list, grid, and pie (similar to circular interface). They found that the users tended to click at the top and bottom of these three interfaces. The users preferred pie and grid over list and their confidence during decision-making was highest for the pie interface. In addition to clicks, users' gaze can be used for behavior analysis over item lists. In another study Chen et al. [4] compared different layouts in which the recommended items can be presented: (a) a simple list, and organization interface with (b) vertical layout and (c) quadrant layout. In organization layout items are grouped by the category (category is an explanation of representative properties of the item). They found that users were more likely to buy items grouped by the category titles and also had more fixation on products in this interface.

Users' clicking activity can be biased towards the rank of the items, such that the probability of a click on an item depends on the rank of the item – *position bias* [7, 12]. When position bias occurs, a click might be a result of the position and not of a relevance. This poses an issue namely during the user feedback interpretation where the raw user clicks cannot be considered to be representatives of the actual users' preferences.

In position bias research authors mainly focus on items represented by their textual characteristics. There are several domains where images associated with items play an important role when a user interacts with the interface (such as movies [25], or fashion [24, 27]). Since items are represented by images, they pose a valuable information for users about the items' characteristics.

Visual bias occurs when users' behavior is changed due to the fact that they perceive the items that are represented by the images. Similarly to the position bias problem, when the visual bias happens and we analyze users' behavior, his/her activity can be misinterpreted and we cannot confidentially decide whether the user's click on the item was invoked by the user's preference (to the item) or by the appearance of the item.

Visual attractiveness and saliency of images on the Web have been already well-studied. Nielsen et al. [16] revealed that, for a majority of tasks, users follow an F-shaped pattern while reading the website's textual content. Nielsen's study was further extended by Shrestha and Lenz [20, 21], who studied the user's gaze over image-heavy e-commerce website. Here, the participants were exposed to two types of interfaces – image-based and text-based and were given two basic tasks: browsing and searching. They found out that the users exposed to a page containing images focused mostly on images themselves. Moreover, during both browsing and searching tasks, the participants did not follow the *F-pattern*.

The visual bias in the domain of recommendation is not well understood. In the stereotypical case of vertical one-dimension textual item lists we can rather safely assume the sequences according F-pattern [16, 19], where mainly the position bias plays a role [7, 12]. However, other setups, such as grid-based interfaces, although heavily used in practice, are investigated by few works only ([15, 26]). There is already a knowledge of the shift of user behavior when facing the interfaces containing images, but there is an important open problem – how to utilize this knowledge in an evaluation of the recommender system. Several models have been proposed ([6, 26]) that use gaze data in order to predict user behavior and recommend items. However, current state-of-the-art either cover the textual representation of the items, or a combination of textual and image representation. Thus, there is a need to investigate textual and image representation separately to better understand visual bias and interpret users behavior when interacting with recommended items.

3 STUDY OF VISUAL BIAS

We identified three steps that need to be reflected during the study of the visual bias: detection, quantification, and interpretation of the visual bias (in recommendation process). In this paper, we focused on the detection of visual bias.

For this purpose, we designed and conducted an eye-tracking user study. We focused on the role of visual bias in changes of user behavior, thus we chose two types of interfaces where the items were represented either with text, or images. The study was done in the domain of movies where the images are already a primary representation of the items. The movies were presented to users in a circular layout, where each card in the circle represented one movie.

The research question of the study was as follows: *How the user behavior differs in the task of picking a movie in case if the movies are represented as posters in comparison to text?*.

3.1 Study Scenario

The main task of the study participants was to *select one movie from the presented on the screen.* Movie lists were generated either randomly or based on the preferred genres that participants selected at the beginning of the study. After viewing all of the 24 stimuli (192 different movies), the participants were asked to indicate the movies that they had previously watched (i.e., the movies they watched before participating in the study). The list of the movies to check contained only those movies that were displayed at stimuli.

Dataset. We used the MovieLens 20M [11] dataset containing 27 278 movies. Additional metadata (actors, directors, plot) and posters were obtained from The MovieDB API[1]. We removed the movies without the poster, the most popular movies (TOP 10% of

[1] https://www.themoviedb.org/documentation/api

the movies based on their average rating), and the movies released before 1990 (to reflect the age of the participants).

Presentation of the items. To ensure that the initial probability of the first gaze would not be influenced by previous attention [9], items were positioned approximately into the circular shape (Figure 1). Moreover, participants were asked to fixate on a small target in the middle of the screen for 4 seconds before each set of movies was displayed to them. The target was also used for the validation of calibration as described below.

To compare the behavior of the participants when interacting with the texts and images, two types of items representations were used: (a) text-based that used only the title of the movie, (b) image-based that used only the poster of the movie (Figure 1). Title or poster were placed into the rectangle with fixed dimensions (card). In text-based representation, the title of the movie was vertically centered and there was a rectangle with gray-colored background of a fixed height to ensure that the text would not cause any visual bias. If a participant chose to show the details of the movies (using the button *Detail*), they were displayed within the rectangle.

Stimuli generation. To cover various scenarios of movie selection, we opted for the several approaches of generating movie sets (Table 1). Firstly, we utilized three genres that the participant picked as preferred at the beginning of the study. For each genre we generated one text-based stimulus (row A) and three image-based stimuli (row B). In addition, we randomly selected three non-preferred genres and for each generated one image-based stimulus (row C) to capture possible differences across various genres. We generated 9 stimuli (D, E), where the movies were generated randomly from the whole dataset.

Table 1: Strategies for stimuli generation. We did not use equal number of the stimuli, since our goal was to gather more data for image-based stimuli (to use them for future studies). When directly comparing text-based and image-based strategies, we firstly normalized the results.

	Source of the movies	Presentation	No. of stimuli
A	Preferred genres	Text	3
B	Preferred genres	Image	9
C	Non-preferred genres	Image	3
D	Random	Text	5
E	Random	Image	4

Each movie was presented to a participant in the entire session only once (to prevent any priming and memory effects). Moreover, we randomized the order of movies within stimuli and the order of the stimuli themselves. If two participants picked the same genre, they both encountered the same set of movies at some point in the session (though in a different order).

3.2 Data Collection and Preprocessing

The study was conducted in controlled, eye-tracking laboratory conditions in User eXperience and Interaction Research Center at Faculty of Informatics and Information Technologies, Slovak University of Technology in Bratislava[2] [1]. All the instructions and parts of the study were presented as a website placed within the Tobii Studio[3] test suite. We collected user behavior data in the form of clicking and eye-tracking activity. Eye-tracking data were collected using the Tobii X2-60 eye-trackers (60 Hz sampling freq.) with the screen resolution 1920x1200px. Validations of eye-tracker calibration were performed at the beginning, during the study, and at the end of the study. Average precision (based on [22]) for all the participants was 0.38 degree (median 0.30 degree), average accuracy 0.93 degree (median 0.71 degree).

Figure 1: Example of two stimuli containing the movie cards (left – text-based, right – image-based stimulus). Only one type of representation was shown per stimulus. The circular layout was chosen from the eye-tracking methodological reasons to suppress a potential first fixation bias. After a stimulus was displayed, users had to fixate at the middle cross.

An analysis of the results was based on the eye-tracking data. We converted raw gaze data to fixations using the *I-VT filter* [17]. To identify which items (i.e., movies) attracted the users, we considered the movie cards to be the Areas of interest (AOIs). There were exactly 8 AOIs corresponding to the 8 cards (with unique movies) for each stimulus. The main goal was to identify fixations that matched particular AOIs. We sequentially analyzed each fixation and checked within which AOI it fitted.

The fixations that matched the interaction buttons (*Select*, *Detail*) were omitted. To account for small calibration errors in AOI hit detection, we experimentally set the absolute error tolerance of $\Delta = 5px$ (set based on our observations; larger value of Δ tended to lower the detection accuracy). A fixation was considered as an AOI hit, when it fell within this 5px boundary around the AOI rectangle. To remove the users and stimuli with corrupted eye-tracking data, we calculated the median Euclidean distance between the fixation and the validation point and removed users and stimuli with distances above $100px$ (based on the validation target that has size 100x100px).

4 STUDY RESULTS

There were totally 64 participants in the study (45 males, 19 females); varying in age 15-27 ($\sigma = 3.42$). Most of them were high school and university students, the rest adults of various occupations.

[2] https://uxi.sk/
[3] https://www.tobiipro.com/product-listing/tobii-pro-studio

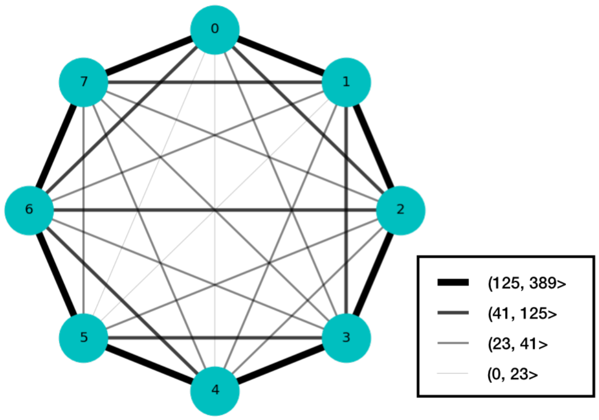

Figure 2: AOIs transition graph of image-based stimuli. Thickness of the edges corresponds to the number of transitions between AOIs (sum over all participants). Intervals are based on quartiles and median of the number of transitions. There is a strong circular pattern of gaze sequences (edges on the perimeter have notably higher number of transitions in comparison to other edges). Since the differences between the text-based and image-based AOI transitions were not notable, we omit the text-based version.

A pilot study was conducted with 3 participants to identify the problems and to adjust the overall design of the experiments. After the removal of error data, we were left with 56 participants and 1,169 stimuli instances to analyze.

4.1 Position Bias

In a circular interface, we hypothesized that the position bias may have an influence on user's behavior in two main ways: (1) there is an AOI that tends to have significantly higher frequency of the first fixations, (2) the order of the visits on the items is similar to the circular order of the items.

Although we used the validation target (in the center of the screen) at the beginning of each stimulus to reset participants' gaze, it was revealed that in most cases they tended to start to fixate mostly in the same position, which is ascribed to the top of the screen. The behavior was not different regardless of the texts or images (based on the Kolmogorov-Smirnov statistic test, stat=0.5, p-value=0.1877).

AOIs transition graph for the images representation is illustrated in Figure 2. It is clear that users tended to follow the circular path while observing the items (visible from the edges on the perimeter with highest number of occurrences). The similar pattern was present in the text type of representation. On the other hand, from the graph, we may reveal that there were also cases, where the users tended to skip some AOIs.

4.2 User Sequence Breaks

To measure the difference in sequences of AOIs visits, the *gaze transition sizes* were calculated for each stimulus. We defined the gaze transition size as the number of AOI that users visited until the next AOI, regardless of the direction (clockwise / counterclockwise). Hence the number of displayed items was 8, the values of gaze transition sizes fit into $\{1, 2, 3, 4\}$. Size of 1 reflects that the user visited the next AOI, sizes above 1 that the sequence of gaze was broken and a participant *skipped* some of the items (a possible indicator of the visual bias in the image-based stimuli).

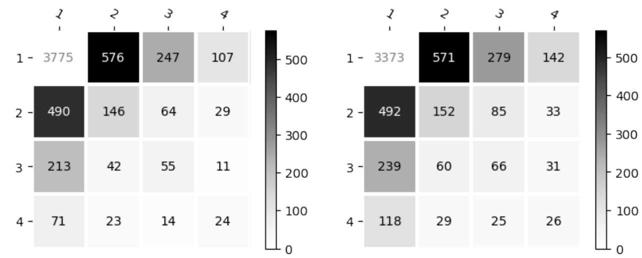

Figure 3: Gaze transition matrices of text-based (left) and image-based (right) stimuli. Rows indices represent the size of the transition in the previous step in a sequence. Column indices represent the size of the transition in the following step in a sequence. Counts are a sum of occurrences over all the participants and all AOIs.

Figure 3 shows the absolute counts of gaze transitions (a sum over all users and all stimuli). We may identify a difference in transition of size 1, where the texts exceed the images. On a counterpart bigger transitions happened more frequently when the user was presented the images. However, in this case the absolute counts may be a bit misleading, since the users tended to behave differently.

To verify whether the behavior varied on a per-user basis, for each user and each stimulus we calculated an average size of gaze transitions, calculated a per-user average, and performed the pairwise t-test between the users (all the populations had normal distribution, based on D'Agostino-Pearson's test, $p > 0.05$). It was revealed that the average size of the gaze transition was significantly different (p-value < 0.01) when comparing text-based (mean=1.28, $\sigma = 0.13$) and image-based (mean=1.36, $\sigma = 0.14$) stimuli. The total count of breaks (i.e., gaze transitions greater than 1) per stimulus was in average also higher for the images (text = 3.49, images = 4.01, $p < 0.01$).

If we examine the image type of representation only, the average *size of the break was largest for the posters of random movie stimuli*. The smaller breaks (size 2) were more frequent for the strategy of preferred genres, whereas the largest breaks (size 4) were more common when the random movies were presented.

5 CONCLUSION

Attention bias poses a fundamental issue when analyzing implicit behavior of the users on the Web. In this paper, we studied two types of attention biases – position bias and visual bias. We proposed and conducted an eye-tracking user study with 64 participants that strives to better understand these biases in the task of selecting an item from the presented recommended items. For the execution of the study, we picked the movie domain.

We compared the users' behavior when the item (movie) was represented either by a text, or by a poster (image) and identified several differences. In both representations of the items a strong position bias occurred. The items were presented in a circular layout to eliminate positon bias. Participants tended to follow its shape while browsing. However, when the participants were exposed to the movie posters, they were more likely to break such a sequence and change their standard gaze path. This behavior was consistent across many participants. Another interesting finding was that even though we instructed the participants to look at the center of the screen before each stimulus, they tended to start their gaze sequence at the AOIs that were placed at the top of the screen.

Our study provides a direct comparison of image-based interface with text-based interface and attempts to quantify the differences. The analysis of user sequence breaks showed that images caused the changes in user behavior from several perspectives. We also found that the user behavior varied based on the shown genre, or a degree of a preference to a genre when the images were presented. Thus, not only the representation might be important, but also the choice of which image to show should be properly examined. We believe that this choice should be also taken into account during the evaluation of user behavior in results of recommendations.

We strive to further investigate user behavior in the recommendation item grids containing images. In our future work, we plan to utilize the measured user behavior and use it in order to take into account both position and visual bias in feedback interpretation task and user behavior modelling.

ACKNOWLEDGMENTS

This work was partially supported by Human Information Behavior in the Digital Space, the Slovak Research and Development Agency under the contract No. APVV-15-0508, and grants No. VG 1/0646/15, No. VG 1/0667/18, and No. KEGA 028STU-4/2017. The authors would also like to thank for financial contribution from the STU Grant scheme for Support of Young Researchers.

REFERENCES

[1] Maria Bielikova, Martin Konopka, Jakub Simko, Robert Moro, Jozef Tvarozek, Patrik Hlavac, and Eduard Kuric. 2018. Eye-tracking en masse: Group user studies, lab infrastructure, and practices. *Journal of Eye Movement Research* 11, 3 (2018). https://doi.org/10.16910/jemr.11.3.6

[2] Rongwei Cen, Yiqun Liu, Min Zhang, Bo Zhou, Liyun Ru, and Shaoping Ma. 2009. Exploring Relevance for Clicks. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management (CIKM '09)*. ACM, New York, NY, USA, 1847–1850. https://doi.org/10.1145/1645953.1646246

[3] Li Chen and Pearl Pu. 2008. A Cross-cultural User Evaluation of Product Recommender Interfaces. In *Proceedings of the 2008 ACM Conference on Recommender Systems (RecSys '08)*. ACM, New York, NY, USA, 75–82. https://doi.org/10.1145/1454008.1454022

[4] Li Chen and Pearl Pu. 2011. Users' Eye Gaze Pattern in Organization-based Recommender Interfaces. In *Proceedings of the 16th International Conference on Intelligent User Interfaces (IUI '11)*. ACM, New York, NY, USA, 311–314. https://doi.org/10.1145/1943403.1943453

[5] Li Chen and Ho Keung Tsoi. 2011. Users' decision behavior in recommender interfaces: Impact of layout design. In *RecSys'11 Workshop on Human Decision Making in Recommender Systems*.

[6] Shiwei Cheng, Xiaojian Liu, Pengyi Yan, Jianbo Zhou, and Shouqian Sun. 2010. Adaptive User Interface of Product Recommendation Based on Eye-tracking. In *Proceedings of the 2010 Workshop on Eye Gaze in Intelligent Human Machine Interaction (EGIHMI '10)*. ACM, New York, NY, USA, 94–101. https://doi.org/10.1145/2002333.2002348

[7] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An Experimental Comparison of Click Position-bias Models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining (WSDM '08)*. ACM, New York, NY, USA, 87–94. https://doi.org/10.1145/1341531.1341545

[8] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of Recommender Algorithms on Top-n Recommendation Tasks. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys '10)*. ACM, New York, NY, USA, 39–46. https://doi.org/10.1145/1864708.1864721

[9] Andrew T. Duchowski. 2007. *Eye Tracking Methodology: Theory and Practice*. Springer-Verlag, Berlin, Heidelberg.

[10] Evgeny Frolov and Ivan Oseledets. 2016. Fifty Shades of Ratings: How to Benefit from a Negative Feedback in Top-N Recommendations Tasks. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, USA, 91–98. https://doi.org/10.1145/2959100.2959170

[11] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. *ACM Trans. Interact. Intell. Syst.* 5, 4, Article 19 (Dec. 2015), 19 pages. https://doi.org/10.1145/2827872

[12] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately Interpreting Clickthrough Data As Implicit Feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '05)*. ACM, New York, NY, USA, 154–161. https://doi.org/10.1145/1076034.1076063

[13] Tilke Judd, Krista Ehinger, Fredo Durand, and Antonio Torralba. 2009. Learning to predict where humans look. In *2009 IEEE 12th International Conference on Computer Vision*. 2106–2113. https://doi.org/10.1109/ICCV.2009.5459462

[14] Michal Kompan, Ondrej Kassak, and Maria Bielikova. 2017. Beyond User Preferences: The Short-Term Behaviour Modelling. In *Proceedings of the RecSys'2017 1st Workshop on Temporal Reasoning in Recommender Systems (CEUR Workshop Proceedings)*, Vol. 1922. CEUR-WS.org, 1–3. http://ceur-ws.org/Vol-1922/paper1.pdf

[15] Wanxuan Lu and Yunde Jia. 2014. An Eye-Tracking Study of User Behavior in Web Image Search. In *PRICAI 2014: Trends in Artificial Intelligence*. Springer International Publishing, Cham, 170–182.

[16] Jakob Nielsen. 2006. F-Shaped Pattern For Reading Web Content, Jakob Nielsen's Alertbox. https://www.nngroup.com/articles/f-shaped-pattern-reading-web-content-discovered/. (2006). [Online; accessed 25 February 2018].

[17] Anneli Olsen. 2012. The Tobii I-VT fixation filter. Algorithm description. https://www.tobiipro.com/siteassets/tobii-pro/learn-and-support/analyze/how-do-we-classify-eye-movements/tobii-pro-i-vt-fixation-filter.pdf. (2012).

[18] Claudia Roda. 2011. Human attention and its implications for human–computer interaction. *Human attention in digital environments* 1 (2011), 11–62.

[19] Andrew Schall and Jennifer Romano Bergstrom. 2014. 1 - Introduction to Eye Tracking. In *Eye Tracking in User Experience Design*, Jennifer Romano Bergstrom and Andrew Jonathan Schall (Eds.). Morgan Kaufmann, Boston, 3 – 26. https://doi.org/10.1016/B978-0-12-408138-3.00001-7

[20] Sav Shrestha and Kelsi Lenz. 2007. Eye gaze patterns while searching vs. browsing a Website. *Usability News* 9, 1 (2007), 1–9. http://usabilitynews.org/eye-gaze-patterns-while-searching-vs-browsing-a-website/

[21] Sav Shrestha, Kelsi Lenz, Barbara Chaparro, and Justin Owens. 2007. "F„ Pattern Scanning of Text and Images in Web Pages. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 51, 18 (2007), 1200–1204. https://doi.org/10.1177/154193120705101831

[22] Tobii Technology. 2011. Eye tracker accuracy and precision. (2011). https://www.tobiipro.com/learn-and-support/learn/eye-tracking-essentials/what-affects-the-accuracy-and-precision-of-an-eye-tracker/ [Online; accessed 04 May 2018].

[23] Licai Wang, Xiangwu Meng, Yujie Zhang, and Yancui Shi. 2010. New Approaches to Mood-based Hybrid Collaborative Filtering. In *Proceedings of the Workshop on Context-Aware Movie Recommendation (CAMRa '10)*. ACM, New York, NY, USA, 28–33. https://doi.org/10.1145/1869652.1869657

[24] Anna Wroblewska and Lukasz Raczkowski. 2016. Visual Recommendation Use Case for an Online Marketplace Platform: Allegro.Pl. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '16)*. ACM, New York, NY, USA, 591–594. https://doi.org/10.1145/2911451.2926722

[25] Lili Zhao, Zhongqi Lu, Sinno Jialin Pan, and Qiang Yang. 2016. Matrix factorization+ for movie recommendation. *IJCAI International Joint Conference on Artificial Intelligence* (2016), 3945–3951.

[26] Qian Zhao, Shuo Chang, F. Maxwell Harper, and Joseph A. Konstan. 2016. Gaze Prediction for Recommender Systems. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, USA, 131–138. https://doi.org/10.1145/2959100.2959150

[27] Zhengzhong Zhou, Yifei Xu, Jingjin Zhou, and Liqing Zhang. 2016. Interactive Image Search for Clothing Recommendation. In *Proceedings of the 2016 ACM on Multimedia Conference (MM '16)*. ACM, New York, NY, USA, 754–756. https://doi.org/10.1145/2964284.2973834