PREFERENCE ELICITATION THROUGH MOUSE CURSOR MOVEMENTS – PRELIMINARY EVIDENCE

Research in Progress

Abstract

Customer preference elicitation is a challenging task with significant practical implications for online shopping. Current methods often put considerable burden on the customers, i.e. through questioning, and could highly benefit from a more accurate estimation of customer weights of product attributes, in particular in an initial purchasing phase. Our goal is to automatically derive attribute weights by recording and analyzing mouse cursor movements. The suitability of the proposed design is confirmed by experimental evaluation, i.e. we found a highly significant correlation between the time people spent investigating a product attribute and their self-reported importance ratings. Our proposed Web page design might also reduce the risk of information overload.

1 Introduction

Online shopping has evolved into a major part of overall economic activity. For example, in the United States, online retail sales totaled more than US$340 billion in 2015, accounting for 7.2 percent of total retail sales (DeNale and Weidenhamer, 2016). However, as online customers lack the ability to physically experience products, customers often rely on product recommendations; for example, McKinsey reported that 35% of Amazon.com's sales are recommended products (MacKenzie et al., 2013). Likewise, especially for complex products, the way detailed product information is presented, is an important factor influencing purchasing decisions (e.g. Franke et al., 2009, Hong et al, 2004, Mandel and Johnson, 2002).

When making purchase decisions, online customers typically follow a two-stage process (Häubl and Trifts, 2000). First, customers consider a large set of candidate products—on one or several Web pages—that they narrow down to a smaller subset of potential alternatives that seem to meet their needs. The customers then assess this subset in more detail through relative comparisons across products based on desirable attributes, before choosing the product that best matches their needs (Ahuja, 2003).

To provide the best (most relevant) recommendations and support the decision process, it is important to understand customers and their preferences (Xiao and Benbasat, 2007). An important part of understanding such preferences is understanding which product attributes customers consider more or less important (i.e., attribute weightings). For example, in multi-attribute decision making it is common to compute the value of an alternative by computing the sum of the product of an attribute multiplied by its weight (Barron and Barrett, 1996; Hwang and Yoon, 1981; Stillwell et al, 1987).

However, in contrast to sales clerks in physical retail stores, who can infer preferences when interacting with a customer or by observing a customer's physical appearance, clothing, or behavior, online retailers often have limited possibilities to infer customers' preferences. Inferring customers' preferences—and providing suitable recommendations—is particularly difficult when no or little information is available, such as when a new customer has not viewed any products before (often referred to as the "cold start problem"; e.g. Bobadilla et al., 2013; Lam et al, 2008; Zigoris and Zhang, 2006).

In this research-in-progress paper, we propose a novel method for implicitly eliciting customers' attribute weightings by using a novel Web site design and investigation of mouse cursor movements. We demonstrate that it is possible to infer attribute weightings by recording and analyzing mouse cursor movements. In the next section, we will briefly discuss preference elicitation before presenting our method, experimental evaluation, and results.

2 Customer Preference Elicitation

Providing relevant recommendations to online customers is challenging as preferences are influenced by many factors (e.g., Holzwarth et al., 2006). Further, people do not necessarily have stable preferences (Holzwarth et al., 2006; Kramer, 2007), and even the choice of included attributes when performing preference elicitation influences a person's preference function (Xiao and Benbasat, 2007). Preferences are often ill-defined, and it has been advocated that people should understand the construction of their preferences (Kramer, 2007).

Preference elicitation is typically based on features, or needs (Xiao and Benbasat, 2007). In feature-based preference elicitation the customer specifies the desired features; this information is used to come up with a reduced set of choices and recommend suitable alternatives. Needs-based preference elicitation requires customers to provide input about their intended usage of the product as well about themselves.

There are two principal methods for eliciting preferences: implicit and explicit. Implicit methods include visual observations of customer interactions (e.g. eye tracking) and asking questions that are not directly related to preferences but still provide information about the customer (and his/her preferences), such as past product purchases. For example, Zdziebko and Sulikowski (2015) tried to predict users' product ratings using a large number of variables obtained from user interaction, including the amount and duration of interaction with a product site such as page visit time, distance covered by the user's mouse cursor, and other variables. Another way of obtaining features and preferences is to extract product features from customer feedback; for example, Popescu and Etzioni (2007) use text mining techniques to identify product features and opinions in online reviews. However, the usefulness of implicitly elicited preferences has been debated (e.g., Zigoris and Zhang, 2006).

Explicit methods involve directly asking customers for their preferences. For example, given a product, customers might have to explicitly state for each product feature, whether they prefer a product with a different value than the shown one (Pu and Chen, 2009). Similarly, (large) online shops such as eBay allow filtering products by specifying values or ranges for attributes. Generally, explicit preference elicitation methods lead to better decision quality but require more effort (due to user input) than implicit methods (Xiao and Benbasat, 2007). Therefore, it seems desirable to have an explicit method with minimal impact on the customer.

There exists a variety of different systems for different purposes that use either implicit or explicit methods, such as travel assistance systems, apartment rental support systems, or (virtual) sales clerks (Chen and Pu, 2004). In practice, to the best of our knowledge large online retailers, such as Amazon, Walmart, and Alibaba seem to prefer non-intrusive, implicit methods in the sense that customers are not questioned explicitly to specify product preferences. Rather, to support the customer in the decision process, online retailers use different forms of recommendation systems (Bobadilla et al., 2013). For example, collaborative filtering is based on products that a user was previously interested in. This set of products is then matched either on a customer-to-customer basis or on an item-to-item basis. Matching products on a customer-to-customer basis involves finding customers that have looked at a similar set of products and recommending the most frequently purchased (or viewed) items of the matched customers, based on the assumption that tastes among different customers are similar. Despite its widespread use, this method suffers from an obvious shortcoming: Two customers might have purchased the same products for different reasons, i.e. due to different preferences. Therefore, basing recommendations only on matching customers' purchases (or interests) may lead to suboptimal recommendations. Matching products on an item-to-item basis relies on comparing attributes of items and recommending similar items (Sarwar et al, 2001). As typically, not all attributes are of the same importance, a weighted distance metric to identify most similar products (or nearest neighbors in the product attribute space) is commonly employed. However, the relevance of particular product attributes is usually not known for a customer, in particular, if he/she has not purchased the product before or is a first time visitor. Though in some cases information about demographics of a customer is available that can partially help to guide recommender systems, it is clear that more detailed information is needed to improve the relevance of the recommendations provided (Bobadilla et al., 2013).

In order to address the shortcomings of current approaches, we seek to develop a non-intrusive way of combining explicit and implicit methods for eliciting preferences by analyzing customers' mouse cursor movements as they interact with a Web site displaying product comparisons. In particular, eye movement studies have indicated that higher fixation frequency can be indicative of greater interest in the target (Jacob and Karn, 2003). Since mouse movements and mouse cursor movements correlate (Chen, Anderson, and Sohn, 2001; Roddon and Fu, 2007), it stands to reason that online customers investigate attributes that are important to them longer than unimportant ones. Based on this assumption, we develop a system that allows eliciting preferences and attribute weightings by analyzing the time users spend with their mouse cursor examining certain product features.

3 Web Page Design and Interaction

Products are generally characterized through a set of attributes that are often depicted in a list or grid (to facilitate product comparisons). These attributes include not only clearly measurable aspects such as weight, price, or customer ratings, but also less tangible aspects such as product design. Commonly, products are presented using a listing of these attributes, listing attribute names and values as shown in Figure 1 for prominent online stores such as Amazon and BestBuy. While allowing for easily comparing products, information overload can be a serious concern, especially for complex products with a large number of attributes (as illustrated in the right part of Figure 1, which shows a part of an attribute-product comparison grid).

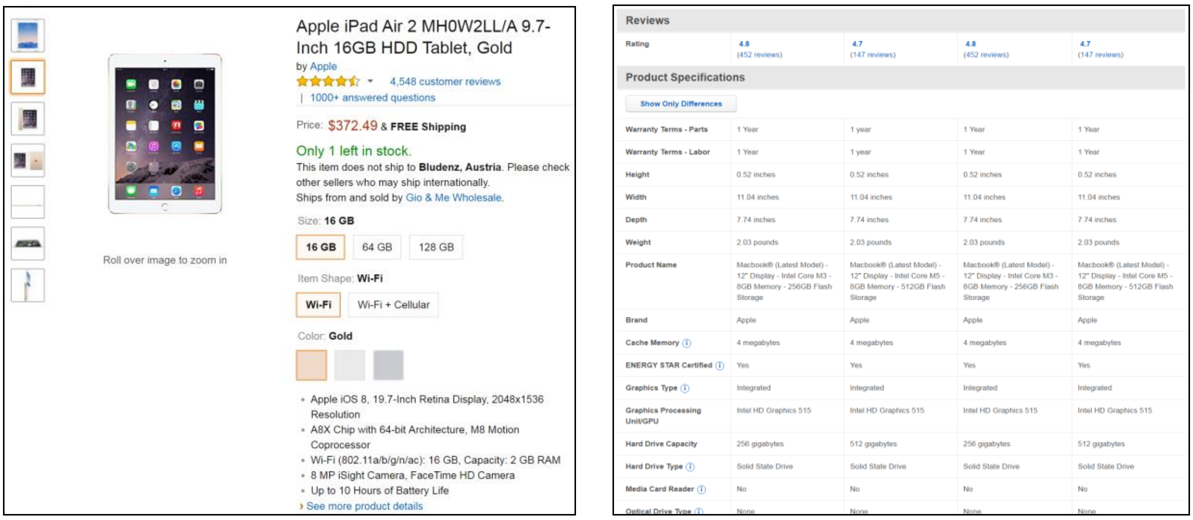

Figure 1: *Left:* Typical structure of a product display within an online store (Amazon.com). *Right:* Attribute-product comparison grid (Bestbuy.com).

We address the problem of inferring customer preferences by using a novel design for comparison grids that only displays product attributes when the user's mouse cursor hovers over the attribute; our new design is combined with a fine-grained data collection and analysis of customer interaction data that goes beyond clicks and time spent on a Web page.[1]

3.1 Web Design: Attribute-Product Visualization and Site Interaction

Our goal was to develop a design that allows reliably obtaining user preferences while minimizing efforts and reducing the risk of information overload. To this end, we modified the traditional attribute-product comparison grid so as to display the attribute names only, and reveal the attribute values when a customer hovers with the mouse cursor over the attribute (see Figure 2). In other words, the attribute values are hidden, and are only revealed when a user explicitly wishes to see the attribute value. This allows us to measure the time users spend examining the values for a particular attribute, and allows us to infer the attribute's relative importance.

Our preference elicitation method is of explicit nature, since a user must move the mouse to see the value of an attribute. This requirement results in a different user experience. In particular, usability might be impaired compared to traditional Web shop designs that would contain all values upfront. Yet reducing the amount of shown information might also improve usability by reducing the potential for information overload arising from the presentation of a large number of attributes. Thus, our design might be particularly favorable to consumers suffering more easily from information overload.

[1] Pages for mobile devices are often designed differently due to reduced screen size and different interaction technologies, i.e., touchscreens rather than pointing devices such as computer mice or trackpads. In this study, we focus on non-mobile devices.

However, employing such a design might also come at the price of usability for some consumers who do not wish to move the pointing device. To some extent, this is expected in an initial phase, since the design is novel and, thus, users are not habituated to it. Still, since most users navigate a mouse pointer frequently and, arguably, with very little effort—both physically and cognitively, it seems reasonable to assume that usability is not strongly affected by the requirement to move the mouse.

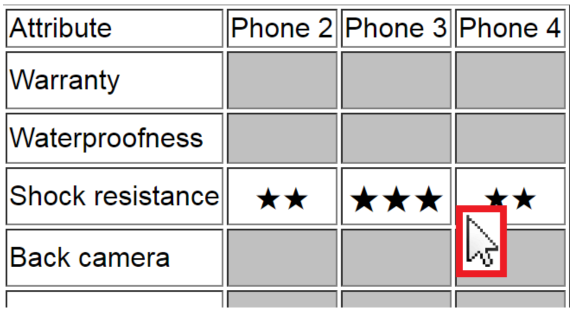

Figure 2: Product comparison grid for three smartphone models. Values are only shown for an attribute on which the mouse cursor is positioned.

3.2 Interaction Logging and Analysis

In this study we focus on non-mobile devices with standard input methods consisting of a keyboard in combination with a pointing device such as a mouse or a trackpad; thus, our focus of analysis and measures lies on measuring interaction using data originating from these devices. As Web shop navigation and product selection can almost completely be done without a keyboard, our emphasis can be further narrowed down to gathering information stemming from the pointing device. This human browser interaction is captured using a log of events, consisting of timestamps, mouse screen coordinates, and the names of the attributes that reside at the specific mouse location.

As indicated, we argue that online customers focus more on attributes that are more important to them, which is in turn reflected in the duration they investigate a specific feature, and seek to derive customer preferences from the collected interaction data. However, customers might not behave rationally, e.g. after multiple decisions some people suffer from decision fatigue (Tierney, 2011) and select a product without proper consideration of their actual desires. A customer might also spend considerable time comparing a small set of attribute values of different products that involve trade-offs, e.g. one product having longer warranty than other products but having less favorable weight. Therefore, the decision process might be more focused on these potentially less important attributes causing more cognitive load that is likely reflected in prolonged attention to these attributes than for more important attributes. Indications are given by eye-tracking studies, where it has been reported that the duration of a fixation is dependent on the processing time used for the fixated object (Just and Carpenter, 1976; Pelz and Canosa, 2001). Thus, we conducted an experiment to determine in how far customers investigate product attributes that they regard as more important longer than product features that they deem as less relevant.

4 Method

To examine the relationship between *mouse hovering time* and *attribute weighting* we have conducted a study using a mock-up page of an online shop and asked participants to select products as well as to weight the importance of distinct attributes.

4.1 Participants

We recruited 50 participants from www.prolific.ac. Forty-four percent were from the United Kingdom, 42% from the United States, and the remaining participants were from various countries in Europe. 80% were native English speakers. 48% were male and 52% female. Mean age was 30.47 years. Though in principle our design might prove effective on handheld devices, in this study we focused on mouse cursor movements. Therefore, we removed data from participants using touch screen devices, leaving us with a total of 46 participants. We used a repeated-measures design, where each participant had to evaluate eleven product attributes in four different scenarios; thus, we ended up with 1932 attribute observations (for the final data analysis, we removed the first scenario, which was used to familiarize the participants with the interface, resulting in a final dataset of 1426 attribute observations.).

4.2 Prototype Implementation

There is a large range of data collection frameworks stemming from industry (e.g. Google Analytics) as well as from academics (e.g. Zdziebko and Sulikowski, 2015). However, freely available tools such as Google Analytics do not provide data (such as mouse cursor movement data) at a sufficient level of detail, and thus limit the possibilities for in-depth analysis. Therefore, we implemented the system ourselves, i.e. a Web page, data collection and analysis. We relied on general purpose Web design and programming tools such as Javascript, HTML, jQuery, Python (and its libraries).

4.3 Procedure

Table 1 provides an overview of the different scenarios. In each scenario, the participants made a purchase decision for either smartphones or cars (with two or three alternatives; see Figure 3 for the smartphone version). For each product, we chose eleven attributes that we found in product descriptions of several online stores. We randomized attribute order as well as product characteristics for each scenario and participant as follows. Attribute order was chosen uniformly at random, the value of attributes of a product was assigned uniformly at random under the constraint that each product obtained the same number of attributes having a one-, two-, and three-star rating. We chose star ratings rather than actual values (such as display sizes, camera resolution, or weight) to reduce the risk of decision fatigue (Tierney, 2011) that can result from having to compare more complex attribute values. The first scenario was intended as training to make users familiar with the interface; thus, we have only used the remaining three scenarios for data analysis. To obtain the attribute-weighting data, we asked participants after Experiment 2 (i.e. smartphones) and after Experiment 4 (i.e. car) to provide a weight for each attribute (Figure 3 shows the three steps for the smartphone experiment).

Table 1: Overview of Scenarios

Scenarios	Product type	Available alternatives	Number of Attributes
1	Smartphone	2	11
2	Smartphone	3	11
3	Car	2	11
4	Car	3	11

4.4 Measures

Attribute weighting. For each attribute, participants were asked to provide a weighting score, ranging from 0 to 100.

Mouse hovering time. We captured a log of mouse events. Each event consisted of a timestamp, mouse screen coordinates, and the name of the HTML element that resides at the specific mouse location. As postprocessing we extracted the total hovering time for each attribute that the mouse cursor resided on the attribute row, i.e. either the attribute name or a value of the attribute held by one of the products.

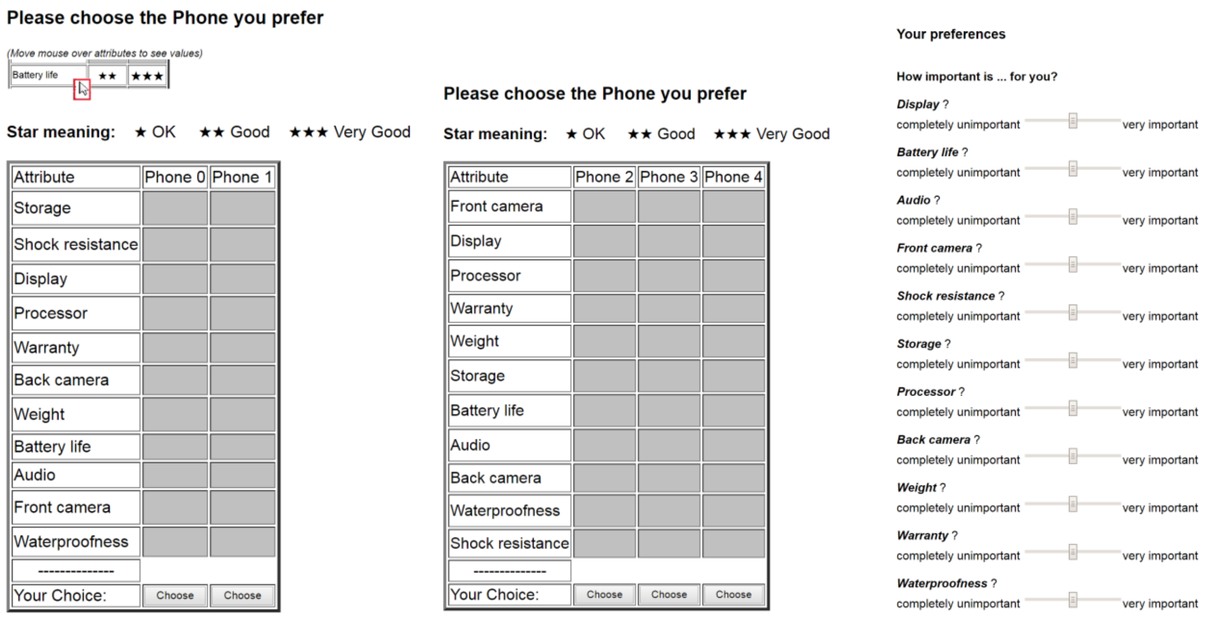

Figure 3: Procedure: *Left:* Scenario 1, *center:* Scenario 2, *right:* attribute weighting

5 Results

To estimate the relationship between mouse hovering time and attribute weighting, we specified a linear mixed-effects regression model. In particular, as observations from the same participant or the same scenario might be correlated, mixed-effects models account for repeated measures (i.e., each participant evaluated eleven attributes) by allowing each participant and each scenario to have a varying intercept in the model. As such, the model accounts for unobserved between-subject heterogeneity (e.g., as the experiment was conducted online, we were not able to observe individual environments) in addition to observed heterogeneity (such as age or gender). We specified the following varying-intercept model:

$$Attribute\ weighting_{i[j]*k} = \alpha_{i[j]*k} + \beta_1 \cdot mouse\ hovering\ time_{i[j]*k} + \gamma' \cdot Controls_{i[j]*k} + u_i + u_k + \varepsilon_{i[j]*k}$$

for $i = 1,\dots, 46$ participants and $j = 1,\dots, 11$ attributes, and $k = 1,\dots, 3$ scenarios. The subscript $i[j]*k$ indexes an attribute j, which is clustered within a participant i (repeated-measures design), where k is the number of scenarios. $\alpha_{i[j]*k}$ represents the individual intercept, β_1 is the effect of mouse hovering time, $Controls_{i[j]*k}$ are the control variables Age and Gender. u_i and u_k are random effects designed to capture the correlation between 1) attribute weights j from the same participant i, and *2)* attribute weights j from the same scenario k.

Table 2 shows the estimated coefficients, standard errors, p-values, and confidence intervals of the fixed effects as well as the variances of the random effects. The results indicate a significant positive relationship between *mouse hovering time* and *attribute weighting* ($\beta = 2.38$; $p < 0.01$; Model 1). The effect was robust when adding control variables ($\beta = 2.39$; $p < 0.01$; Model 2).

Table 2: Mixed-effects regression models

Dependent variable: Attribute weighting						
	Model 1: without controls			**Model 2: with controls**		
Fixed effects	β	p	CI	β	p	CI
Mouse hovering time	2.38 (0.36)	***	1.67 – 3.09	2.39 (0.37)	***	1.66 – 3.12
Age				0.31 (0.12)	**	0.07 – 0.55
Gender				5.43 (2.95)		-0.36 – 11.22
Intercept	62.37 (1.68)	***	59.08 – 65.65	50.09 (4.69)	***	40.91 – 59.27
Random effects						
σ^2	535.96			509.97		
σ^2_1 (Participant; N = 46)	84.07			70.51		
σ^2_2 (Experiment; N = 3)	0.96			1.60		
AIC	13,100.6			12,180.9		
BIC	13,126.9			12,217.2		
Observations	1426			1333		

*Notes: * p < .05 ** p < .01 *** p < .001; standard errors are in parentheses; we had some missing values in Model 2 (for Age and Gender), thus, the number of observations is slightly different from Model 1.*

6 Discussion and Conclusion

In this research-in-progress paper we have proposed a novel design for product comparisons. We demonstrated that by only revealing attribute values when a user's mouse cursor is placed over an item and analyzing users' mouse cursor movements, it is possible to infer the importance a user places on an attribute (i.e., attribute weights). Therefore, we have contributed to existing research on customer preference elicitation by proposing a novel method for obtaining customer weights in an online shopping scenario that comes with little effort for the customer. From a practical standpoint, this novel design can help to provide more relevant recommendations by providing information about user's preferences; at the same time, this design can help to reduce information overload, especially for products with multiple attributes. As such, our research has the potential to have important practical implications for online retailers. However, our current design is also somewhat artificial; therefore, we intend to investigate in more detail questions related to privacy, information overload and usability. Furthermore, we have focused on monitoring user behavior based on mice and trackpads, excluding touchscreens. We also plan to develop a design and conduct a study targeted to users of mobile devices with touchscreens. Likewise, we have thus far only tested our design with a relatively small sample using simulated scenarios. In future studies, we intend to replace these scenarios and mock-up designs with conditions involving actual purchases in a real shop to assess the ability to obtain better-tailored product recommendations.

References

Ahuja, M., Gupta, B. and Raman, P. (2003). "An Empirical Investigation of Online Consumer Purchasing Behavior." In: *Communications of the ACM*, *46*(12), 145-151

Barron, F.H. and Barrett, B.E. (1996). "Decision Quality Using Ranked Attribute Weights." In: *Management Science*, *42*(11), 1515-1523.

Bobadilla, J., Ortega, F., Hernando, A. and Gutiérrez, A. (2013). "Recommender Systems Survey." In: *Knowledge-Based Systems, 46*, 109-132.

Chen, M.C., Anderson, J.R. and Sohn, M.H. (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* (Pp. 281-282). ACM.

Chen, L. and Pu, P. (2004). *Survey of Preference Elicitation Methods*. EPFL-REPORT-52659.

Claypool, M., Le, P., Wased, M. and Brown, D., 2001, January. "Implicit Interest Indicators." In: *Proceedings of the 6th International Conference on Intelligent User Interfaces* (Pp. 33-40). ACM.

Coupey, E., Irwin, J.R. and Payne, J.W. (1998). "Product Category Familiarity and Preference Construction." In: *Journal of Consumer Research*, *24*(4), 459-468.

Denale, R. and Weidenhamer, D. (2016). *Quarterly Retail E-Commerce Sales 1st Quarter 2016*, U.S. Census Bureau News. URL: http://www2.census.gov/retail/releases/historical/ecomm/16q1.pdf (visited on 12/3/2016).

Franke, N., Keinz, P. and Steger, C.J. (2009). "Testing the Value of Customization: When Do Customers Really Prefer Products Tailored to Their Preferences?" In: *Journal of Marketing*, *73*(5), 103-121.

Guo, Q. and Agichtein, E. (2010). "Ready to Buy or Just Browsing?: Detecting Web Searcher Goals from Interaction Data." In: *Proceedings of the 33rd international ACM SIGIR Conference on Research and Development in information Retrieval* (Pp. 130-137). ACM.

Jacob, R.J.K. and Karn, K.S. (2003). "Eye Tracking in Human-Computer Interaction and Usability Research: Ready to Deliver the Promises." In: *The Mind's Eye: Cognitive and Applied Aspects of Eye Movement Research*, Elsevier, Amsterdam, 573-605.

Just, M.A. and Carpenter, P.A. (1976). "Eye Fixations and Cognitive Processes." In: *Cognitive Psychology*, *8*(4), 441-480.

Goldberg, J.H., Stimson, M.J., Lewenstein, M., Scott, N. and Wichansky, A.M. (2002). "Eye Tracking in Web Search Tasks: Design Implications." In: *Proceedings of the 2002 Symposium on Eye Tracking Research & Applications* (Pp. 51-58). ACM.

Hwang, C. L., and Yoon, K. (1981). *Multiple attribute decision making: methods and applications: a state-of-the-art survey*. Berlin, Springer-Verlag.

Häubl, G. and Trifts, V. (2000). "Consumer Decision Making in Online Shopping Environments: The Effects of Interactive Decision Aids." In: *Marketing Science*, *19*(1), 4-21.

Holzwarth, M., Janiszewski, C. and Neumann, M.M. (2006). "The Influence of Avatars on Online Consumer Shopping Behavior." *Journal of Marketing*, *70*(4), 19-36.

Hong, W., Thong, J.Y. and Tam, K.Y. (2004). "The Effects of Information Format and Shopping Task on Consumers' online Shopping Behavior: A Cognitive Fit Perspective." In: *Journal of Management Information Systems*, *21*(3), 149-184.

Kramer, T. (2007). "The Effect of Measurement Task Transparency on Preference Construction and Evaluations of Personalized Recommendations." *Journal of Marketing Research*, *44*(2), 224-233.

Lam, X.N., Vu, T., Le, T.D. and Duong, A.D. (2008). "Addressing Cold-Start Problem in Recommendation Systems." In: *Proceedings of the 2nd international Conference on Ubiquitous Information Management and Communication* (Pp. 208-211). ACM.

Mackenzie, I., Meyer, C. and Noble, S. (2013), *How Retailers Can Keep Up with Consumers.* URL: http://www.Mckinsey.Com/industries/Retail/Our-insights/How-Retailers-Can-Keep-Up-With-Consumers (visited on 11/25/2016).

Mandel, N. and Johnson, E.J. (2002). “When Web Pages Influence Choice: Effects of Visual Primes on Experts and Novices.” In: *Journal of Consumer Research*, *29*(2), 235-245.

Moe, W.W. (2003. “Buying, Searching, or Browsing: Differentiating Between Online Shoppers Using In-Store Navigational Clickstream.” In: *Journal of Consumer Psychology*, *13*(1), 29-39.

Pelz, J.B. and Canosa, R. (2001). “Oculomotor Behavior and Perceptual Strategies in Complex Tasks..” In: *Vision Research*, Vol. 41 Nos 25-26, 3587-3596.

Popescu, A.M. and Etzioni, O. (2007). “Extracting Product Features and Opinions From Reviews.” In: *Natural Language Processing and Text Mining* (Pp. 9-28). Springer London.

Pu, P. and Chen, L. (2009). “User-involved Preference Elicitation for Product Search and Recommender Systems.” In: *AI Magazine*, *29*(4), 93.

Rodden, K. and Fu, X. (2007). “Exploring How Mouse Movements Relate to Eye Movements on Web Search Results Pages.” In: *Web information Seeking and interaction*, 29-32.

Sarwar, B., Karypis, G., Konstan, J. and Riedl, J. (2001). “Item-Based Collaborative Filtering Recommendation Algorithms.” In: *Proceedings of the 10th International Conference on World Wide Web* (Pp. 285-295). ACM.

Stillwell, W.G., Von Winterfeldt, D. and John, R.S. (1987). “Comparing Hierarchical and Nonhierarchical Weighting Methods for Eliciting Multiattribute Value Models.” In: *Management Science*, *33*(4), 442-450.

Tierney, J., 2011. *Do You Suffer from Decision Fatigue?* The New York Times, 21st of August 2011. URL: http://www.nytimes.com/2011/08/21/magazine/do-you-suffer-from-decision-fatigue.html (visited on 11/25/2016).

Xiao, B. and Benbasat, I. (2007). “E-Commerce Product Recommendation Agents: Use, Characteristics, and Impact.” In: *MIS Quarterly*, *31*(1), 137-209.

Zdziebko, T. and Sulikowski, P. (2015). “Monitoring Human Website interactions for Online Stores.” In: *New Contributions in Information Systems and Technologies* (Pp. 375-384). Springer international Publishing.

Zigoris, P. and Zhang, Y. (2006). “Bayesian Adaptive User Profiling with Explicit & Implicit Feedback.” In: *Proceedings of the 15th ACM international Conference on Information and Knowledge Management* (Pp. 397-404). ACM.