Interpreting and Leveraging Browser Interaction for Exploratory Search Tasks

Dominic Stange[2], Michael Kotzyba[1], Stefan Langer[1] and Andreas Nürnberger[1]
[1]DKE Group, Faculty of Computer Science, University of Magdeburg, Germany
[2]Volkswagen AG, Germany

Keywords: Web Usage Mining, Interaction Log Analysis, Exploratory Search, Recommender Systems, Interactive Information Retrieval, Machine Learning.

Abstract: In this paper we introduce a novel approach for modeling and interpreting search behavior for exploratory search by using a so called exploration graph. We use an existing methodology of logging and analyzing user interactions with a web browser and add an additional interpretation step that can be used, e. g. to integrate sensemaking or browsing patterns into the log data. We conducted a user study and are able to show that: (a) interaction logs can be interpreted semantically, (b) semantic interpretations lead to a more connected exploration graph, and (c) multiple (even contradicting) interpretations of the same search behavior may exist at the same time. We also show how our theoretical model can be applied in the area of professional search by incorporating insights gained from the model into novel recommendation and machine learning approaches.

1 INTRODUCTION

In contrast to simple fact-finding search, in exploratory search techniques of learning and investigating are used (Marchionini, 2006). Furthermore, exploratory search is characterized as open-ended and multifaceted with unclear goals (White and Roth, 2009; Wildemuth and Freund, 2012). These characteristics make it difficult to track and examine a user's learning process during search just by considering the interactions with a browser. Especially, in the World Wide Web typical interactions of exploratory search are likely to involve "many impasses, illstructured goals and tasks, navigation and exploration, and substantial influences from the content that is encountered" (Card et al., 2001). Methods that enable efficient support during interactive (exploratory) search could strongly improve the efficiency of the search process, especially in the area of professional search (e.g. technology scouting), which is a common task in large enterprises (Nürnberger et al., 2015). However, one major problem is that an appropriate data model for exploratory search which also enables mechanisms for effective storage and analysis is still missing. Therefore, we propose in this paper a data structure and preliminary results of a small user study towards improved ranking and recommendation based on semantic interpretations of exploration graphs.

This paper is structured as follows: The next section describes related work. Sect. 3 provides a formal model for exploratory search behavior. In Sect. 4 we outline how the model allows us to interpret search behavior and provide a problem definition. Afterwards, the conducted user study and preliminary results are described in Sect. 5. Then, Sect. 6 shows how data about interpreted search behavior can be used as input to a recommender system in a business environment. In Sect. 7 a short discussion and conclusion is given.

2 RELATED WORK

The analysis of search behavior has a long history in research and is rooted in the area of library and information science with the goal to gain a holistic view on the search process and to derive generalized models (Bates, 1989; Marchionini, 1997). A good overview on this macroscopic view is available in (Knight and Spink, 2008). Further research in information retrieval focuses on the user's concrete interactions to build models which represent the user's state or search success. For example, Hassan et al. (Hassan et al., 2010) propose a Markov Model with transition times and show that this approach performs significantly more accurate than traditional relevance-

Stange, D., Kotzyba, M., Langer, S. and Nürnberger, A.
Interpreting and Leveraging Browser Interaction for Exploratory Search Tasks.
DOI: 10.5220/0006372901910197
In *Proceedings of the 19th International Conference on Enterprise Information Systems (ICEIS 2017) - Volume 3*, pages 191-197
ISBN: 978-989-758-249-3

based models for predicting user search goal success. Ageev et al. (Ageev et al., 2011) apply different model approaches to predict user's success for fact-finding search tasks. Card et al. (Card et al., 2001) as well as White and Drucker (White and Drucker, 2007) use browser interaction analysis to create behavioral graphs that model search behavior. Their studies rely on browser interaction logs which store data about how searchers interact with their Web browser. In particular, White and Drucker investigate browser trails, i.e., the sequence of visited websites. For example, if a user visits two websites, a and b, in a row and returns to a to browse to a website, c, the corresponding search trail can be written as $a \to b \to a \to c$. This approach allows to, e. g. analyze and investigate searcher profiles.

Our approach facilitates browser interaction logs, as well. We use them to model a user's search behavior and keep more information about the actual browsing activities in a search than with behavioral graphs. With this model we are also able to create different *interpretations* of search behavior, based on additional assumptions about a user's sensemaking activities during search. The problem of creating such interpretations is formulated in the course of the next sections.

3 A MODEL FOR EXPLORATORY SEARCH

This section outlines a model to formally describe search behavior based on browser interaction. The model is particularly designed with an exploratory search task in mind.

When a person explores the World Wide Web using a Web browser there is always one website visible in the browser. If the URL to this website is, e. g. *www.domain.com* we say that the person is in state ON("*www.domain.com*"). If the person moves to another website by interacting with the Web browser, e. g. by clicking on a hyperlink within the website and the URL to this second website is *www.anotherdomain.com*, the action can be written as $GOTO$("*www.anotherdomain.com*"). Hence, the corresponding search process of the person starting on the *home screen* can be written as an alternating sequence of states and actions:

State 1: ON ("*home screen*")
Action 1: $GOTO$ ("*www.domain.com*")
State 2: ON ("*www.domain.com*")
Action 2: $GOTO$ ("*www.anotherdomain.com*")
State 3: ON ("*www.anotherdomain.com*")
Action 3: ...

We call this sequence of states and actions the person's **search history**. The search history is the result of observing the person's search behavior. A search history contains the websites a person visits during a search process and the order in which they are visited over time. It also contains the actions through which the websites are opened. Note that it is possible for the person to visit a website multiple times during search, e. g. when switching browser tabs, so different states in the search history can have the same website. One could say that through browser interaction the person transitions between different states, where each state is determined by the visited website. These transitions can be written as

$$RESULT(ON(w), GOTO(w')) = ON(w'),$$

i. e. the person moves from state $ON(w)$, where website w is visited, to another state, $ON(w')$, where website w' is visited by performing the $GOTO$-action.

Web browsers typically process the data about these websites and also the browser tabs in which they are shown. Since we want to model a user's interaction with a Web browser our model has to consider all browser tabs and the websites they display. To simplify the model we assume that only one browser window is used by the user. Therefore, only one website is visible in the web browser at any point in time while the rest of the websites are hidden. We call this visible website the **active website** of the browser. The **browser state** can then be defined in terms of the active and all opened websites.

Definition 1 (Browser State)**.** *Given a set $\mathcal{W}$ of websites, a browser state s can be written as*

$$s(\mathcal{W}) = ON(\mathcal{W}) = ON(w_1, w_2, ..., \underline{w_k}, ..., w_{n-1}, w_n),$$

where each website $w_i \in \mathcal{W}$ is shown in a browser tab and the underlined website w_k is the active website of browser state s, with $1 \leq k \leq n$.

At the beginning the browser state is typically empty, i. e. the set $\mathcal{W}$ is empty. We write this as $s = ON(\varnothing)$. Similarly to the user's search history the interactions with the browser result in an alternating sequence of browser states and user actions. We call this sequence the **browser history**.

Definition 2 (Browser History)**.** *Given a set $\mathcal{S}$ of browser states and a set $\mathcal{A}$ of user actions, a browser history h is an alternating sequence*

$$h(\mathcal{S}, \mathcal{A}) = \langle\ s_0,\ a_1,\ s_1,\ a_2,\ \ldots, s_{n-1},\ a_n,\ s_n\ \rangle$$

of browser states $s_i \in \mathcal{S}$ and user actions $a_j \in \mathcal{A}$.

An example of a browser history in which the websites w_1, w_2, and w_3 are visited is:

$$\begin{array}{llll} \langle & s_0 : ON(\varnothing), & a_1 : GOTO(w_1), & \\ & s_1 : ON(\underline{w_1}), & a_2 : GOTO(w_2), & \\ & s_2 : ON(\underline{w_2}), & a_3 : GOTO(w_3), & \\ & s_3 : ON(w_2, \underline{w_3}) & & \rangle. \end{array}$$

Initially, the browser history is empty. Then, websites are opened in the course of the search process. Notice that in the example the third website, w_3, is opened in a new browser tab. Thus, the previous website, w_2, is not removed from the set of websites $\mathcal{W}$ in state s_3. That is, if we only consider the *GOTO*-action it is not possible to "predict" its exact outcome for a browser state. In other words, the *RESULT*-function is not determined for a browser state and the *GOTO*-action. Therefore, we need to observe other user actions to solve this ambiguity. Preliminary analysis shows that the eight actions of Table 1 are well-suited for our search model to work as intended:

Table 1: User actions and their descriptions. The actions are part of our exploratory search behavior model.

Action	Description
URL	open a website in the active browser tab
URL^+	open a website in a new browser tab
CLICK	enter a URL in the address field of the active browser tab
$CLICK^+$	enter a URL in the address field of a new browser tab
CLOSE	close an existing browser tab
SWITCH	switch from one browser tab to another
BW/FW	move backward/forward within the local history of a browser tab
QUERY	enter a search query in the search box of the active browser tab

The actions have their expected effect on a browser state. The *URL*-action opens a website for a given URL which is entered in the address field of the active browser tab. The action removes the active website and adds the new website to the next browser state. For example, if the browser is in state $ON(\underline{w})$ and the user performs the *URL*-action the transition is $RESULT(ON(\underline{w}), URL(w')) = ON(\underline{w}')$. The symbol "+" in the URL^+-action indicates that the website for the URL is opened in a new browser tab. For example, if the browser is in state $ON(\varnothing)$ and the user performs the URL^+-action the corresponding transition is $RESULT(ON(\varnothing), URL(w)) = ON(\underline{w})$. The *CLICK*-action removes the active website and adds the website of the clicked hyperlink to the next browser state. Consistently, the $CLICK^+$-action opens the website of the hyperlink in a new browser tab and, therefore, adds this website to the next browser state without removing the active website. The *CLOSE*-action removes the active website from the next browser state. The *SWITCH*-action changes the active website from one website to another website. The FW/BW-action uses the local history of a browser tab to navigate forward and backward, replacing the active website with the one of the local history. Lastly, the *QUERY*-action issues a search query to the active website, removing it from the next state and adding the corresponding search engine results page (SERP). Using these user actions the browser history for the previous example can be written as:

$$\begin{array}{llll} \langle & s_0 : ON(\varnothing), & a_1 : URL^+(w_1), & \\ & s_2 : ON(\underline{w_1}), & a_3 : CLICK(w_2), & \\ & s_4 : ON(\underline{w_2}), & a_5 : CLICK^+(w_3), & \\ & s_5 : ON(w_2, \underline{w_3}) & & \rangle. \end{array}$$

As can be seen, this browser history describes the person's search behavior unambiguously. Note that sometimes an action is not *applicable* for a browser state. For example, the *QUERY*-action is only applicable if there is a search box on the active website. The actions are atomic in the sense that any of them creates a new browser state. Some Web browsers provide functions that perform more than one atomic action in one interaction, e. g. opening a hyperlink in a background browser tab.

4 INTERPRETING BROWSER HISTORIES

With the help of the eight user actions listed in Table 1 of the previous section we can create a sufficiently rich model of any person's exploratory search behavior. Now, based on the model we want to introduce the noval approach of *interpreting* this behavior formally and also propose the concept of *semantics* in this interpretation.

A structure that takes both the browser states and transitions between them into account is a graph. We call this the **exploration graph**.

Definition 3 (Exploration Graph)**.** *An exploration graph* $x(\mathcal{V}, \mathcal{E})$ *is a cyclic, directed graph with a set* $\mathcal{V}$ *of vertices and a set* $\mathcal{E}$ *of edges, where elements of* $\mathcal{E}$ *are ordered pairs* (v_x, v_y) *of distinct vertices* $v_x, v_y \in \mathcal{V}$.

An exploration graph describes how a user explores an information space based on the websites being visited during search. Every vertex in x represents exactly one website. An example of an exploration graph with five vertices and four edges is shown in

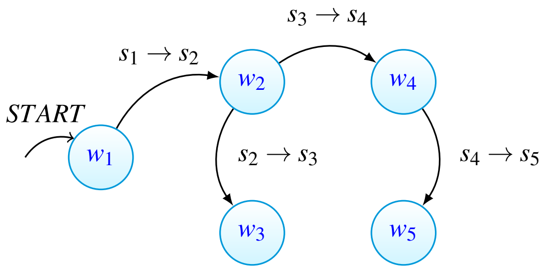

Figure 1: An exploration graph for a random browser history. It consists of five nodes (w_1 to w_5) that represent the visited websites. The edges between the nodes indicate how the information space is traversed. The edge labels indicate between which states the connection is added.

Figure 1. The figure shows an extension to the example from the previous section, where two additional websites, w_4 and w_5, are opened.

An exploration graph visualizes the different exploration *paths* of a user, e. g. $w_1 \rightarrow w_2$ or $w_2 \rightarrow w_4 \rightarrow w_5$. Note the difference of the paths to the one-dimensional strictly chronological sequence $w_1 \rightarrow w_2 \rightarrow w_3 \rightarrow w_4 \rightarrow w_5$ or a browser trail $w_1 \rightarrow w_2 \rightarrow w_3 \rightarrow w_2 \rightarrow w_4 \rightarrow w_5$. The direction of an edge between two websites is determined by the direction of the hyperlink between them, e. g. $w_1 \rightarrow w_2$. For ease of understanding the edges of the example exploration graph are labeled such that they show between which states they are added. In general, it is possible that a user visits the same website multiple times (e. g. following links on different websites to the same target) and, hence, the graph can contain cycles. New paths are created whenever a user enters a URL manually or issues a search query to a search engine. That is, in contrast to browser trails, for example, an edge is not automatically appended to the last active vertex, reducing possible randomness (e. g. the link $w_3 \rightarrow w_2$) and also avoiding heuristics to close a path, c.f. (White and Drucker, 2007).

However, if a user issues a search query based on something she has learned on a website, it is reasonable to assume that this search query continues the path that has led to the website. That is, search queries may refer to the content of visited websites. This relationship is typically not present in the browser history because there is no interaction to be logged. Yet, an edge connecting the website with the SERP of the query can be weighted by the degree/probability that they are *semantically related*. Thus, edges in exploration graphs can express semantics in search behavior. For instance, the probability could indicate how likely it is that the user has used insights gained from a specific website to issue the search query. This is likely to result in a more connected exploration graph, because in addition to the actual browsing the semantic relationships are taken into account. Following this approach also means that there is not a single "true" exploration graph but several possible instantiations that are all different *interpretations* of the same search behavior. In the example shown above not all actions have been interpreted in a way that they create an edge between two websites. For example, a *SWITCH*-action is not represented by an edge. This is in contrast with the interpretation corresponding to the browser trail, which - using our terminology - can be thought of as being based on a search history rather than a browser history. In some cases it may be reasonable to use one interpretation over another depending on the research scenario. There are certainly more types of interpretations possible. A detailed analysis of possible interpretations is beyond the scope of this paper but subject for future work.

The problem of creating an interpretation for (exploratory) search behavior can be stated as follows:
Problem Statement: *Given a browser history $h(\mathcal{S},\mathcal{A})$, find an interpretation function*

$$f : (h) \xrightarrow{theory} x$$

that best represents *the search behavior modeled in h as an exploration graph $x(\mathcal{V},\mathcal{E})$.*

By adding the component *theory* to the problem statement we emphasize that interpretations can make assumptions about a user's search behavior. These assumptions can be used, e. g. to incorporate semantic relationships into the exploration graph. An interpretation can be understood as a projection of observational and/or sensemaking patterns on the data collected via browser interactions of a user.

5 USER STUDY AND PRELIMINARY RESULTS

One goal of the user study is to show how a semantic interpretation of search behavior looks like and how it compares to an interpretation without assumptions. Using insights gained from a pre-study in which we investigated the complexity of exploration graphs, we designed the main study as follows.

Study Design. Participants are asked to search for information about two predefined topics. The topics are selected to satisfy the attributes described in (Wildemuth and Freund, 2012): i. e. learning and investigation as goals; general, ill-structured and open-ended problems; and an involved uncertainty.

One search topic is about the procurement of solar energy panels for a private houshold as adapted to (Gwizdka and Lopatovska, 2009). The other is about body scanners and their affect on man's health

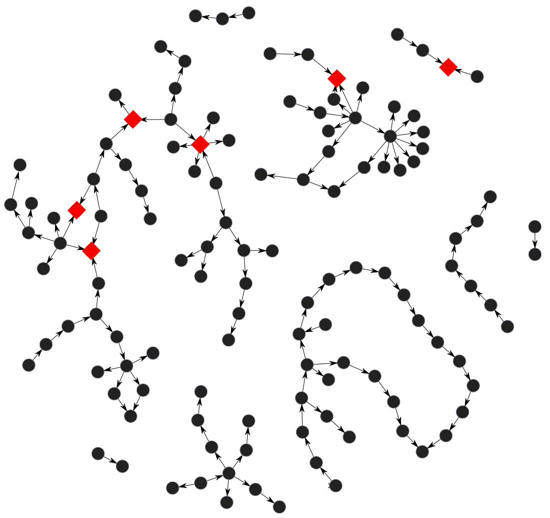

Figure 2: An interpretation of browser interaction data without assumptions about a user's search behavior.

as adapted to (Wu et al., 2012). Participants are advised that each search task should take at least 15 minutes.

Data Preparation. For each search task we create a single browser history by merging the interaction data of all participants into a single log file. Combining the search activity of the participants allows us to "simulate" a more extensive exploratory search with multiple starting points, at least one for each participant. However, we lose the ability to examine individual search profiles, because the resulting interaction log exhibits characteristics from multiple searchers. We then create exploration graphs based on the log files using two different interpretations, one with and one without assumptions about the search behavior.

Preliminary Results. Nine participants aged between 26 and 33 participated the preliminary study. Figure 2 shows the result of an interpretation for the first search topic without assumptions about the search behavior, i. e. only the actual navigations between visited websites are considered in the interpretation. As can be seen, several disconnected graph components represent the search behavior. Most of the starting points are created by visiting general-purpose search engines, like Google. Note that some participants visited the same websites (red rectangles), so even based on this interpretation some of the individual explorations are connected.

A semantic interpretation of the same data is shown in Figure 3. The graph is more connected, because some search engine results pages are connected to relevant websites visited before. To calculate the relevancy of a query we index every website at the time it is visited in a local search engine. Whenever a user issues a query we make a lookup to find relevant websites in this index for the same query terms. At the beginning the index is empty. Only the most relevant (already visited) website is selected to be continued. In this interpretation it is necessary to add a parameter to the interpretation function to set a minimum relevancy threshold. The example interpretation uses no relevancy threshold. A sample of that exploration graph for the first user study is shown in Figure 4. The graph shows an episode of search activity and it outlines how *pre-query browsing* (as opposed to post-query browsing) and search queries can contribute to the same line of thought during a search process.

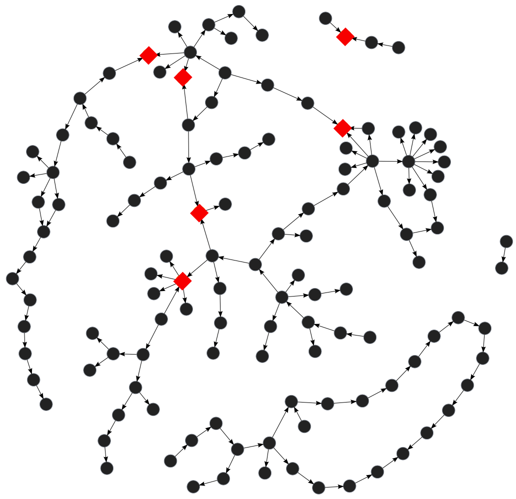

Figure 3: A semantic interpretation of browser interaction data incorporating sensemaking patterns during search.

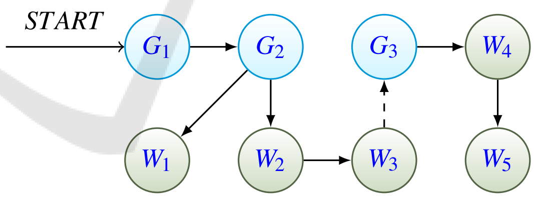

Figure 4: A sample exploration graph for the user study using a semantic interpretation. It contains three Google results pages (G_1 to G_3) and five content websites (W_1 to W_5). The query terms of G_2 are "solar panel advantages over conventional methods". The subsequently visited websites explain pros and cons of solar panels. Website W_3, however, deals with photovoltaics. The query terms of G_3 are "photovoltaic vs solar heat". The websites W_4 and W_5 explain the difference between solar and photovoltaic systems. Due to the used interpretation function the path $G_3 \rightarrow W_4 \rightarrow W_5$ is appended to W_3, indicated by the dotted arrow.

It is also possible to interpret the log data such that search queries are appended to the top k relevant websites at the same time, where k becomes another parameter for the interpretation. This may express uncertainty towards the user's line of thought during exploration. By appending the query website to multiple paths it is possible to represent concurrent interpretations that may contradict each other. For example, a query may be similarly relevant for websites in different paths. Keeping (some of) these interpretations may support a later graph analysis or mining, because more data about the search behavior is available.

6 RECOMMENDATIONS BASED ON EXPLORATION GRAPHS

In the introductory section of this paper we briefly summarized exploratory search as open-ended and multifaceted with unclear goals. Researches often argue that users in exploratory search often do not know exactly what search queries to use in order to find the answers they desire. In this section we want to show how exploration graphs can be leveraged in a recommender system to suggest websites to be visited in such situations. We believe that the analysis of exploration graphs is an exciting approach to develop new support mechanisms for enterprise search systems.

We have implemented a web-based search platform which is currently rolled out in a large automotive company as part of a long-term research project. The overall goal of the project is to make complex search tasks more efficient and support collaboration within teams. Our platform uses a client-side data collection approach. With the help of a browser addon we detect the search actions described in Sect. 3 and send them to a server where they are interpreted as an exploration graph (see Sect. 4) and stored in a graph database. The server processes the stream of interaction data online. We create a single exploration graph for all users. For the recommendations we perform graph mining to find patterns in this exploration graph. Domain experts can use a special-tailored search user interface to query for relevant websites. They are presented with *relevant* websites containing the query terms and *related* websites (which might not contain the query terms) based on the structure and semantics of the exploration graph. This kind of post-query recommendation is a novel approach and particularly suited in business environments where searchers are likely to address similar topics repeatedly over time, e. g. in patent search or technology scouting. Especially, in a scenario where multiple domain experts are interested in the same (or similar) topics, they can benefit from each others' search efforts.

In order to further improve the quality of recommendations we have added an additional action to the list of user interactions: the *LIKE*-action marks the active website as an interesting resource for the current search. So, whenever a user issues a query and browses to a website with relevant content the user can "like" this website. The system then determines the path in the exploration graph from a possible query to this website (also taking into account the length of this path) and recommends it in future searches when similar queries are used. Examples for patterns we examine for our recommendations are:

- Given a website w, find all websites $\mathcal{W}^*$ in the neighborhood of w that are "liked" by at least one member of the group.
- Given a search query q, find those search engine results pages $\mathcal{W}^+$ that are visited based on similar search queries, and identify all websites $\mathcal{W}^*$ in the neighborhoods of all websites $w \in \mathcal{W}^+$ that are "liked" by at least one member of the group.

Next, we want to briefly outline two research directions we want to look into in the future.

Learning to Rank. It is possible to create machine learning algorithms which modify the relevancy of any visited website for a given search query based on the exhibited search behavior using the data of an exploration graph. The *LIKE*-action allows us to use both unsupervised as well as supervised and hybrid approaches. It is even possible to extract additional features from websites, e. g. website content, title, or meta information, and incorporate this into the learning process. Then, suggestions are made based on complex data, taking into account the structural and semantical relationships of observed search behavior.

Learning to Automate. One of the main motivations to initiate the research project with the automotive partner is the long-term goal to (partially) automate complex search tasks. Technologies like focussed crawlers have shown since two decades how automated search algorithms can reduce the effort to perform complex search manually. We believe that using the data provided by exploration graphs we can further improve the results of automated crawlers. The advantage of training such a crawler based on real exploration graphs is that such an approach is flexible with respect to the search topic and the individual traits of the searchers. For example, a domain expert can perform a complex search task manually to provide an initial starting point for the learning algorithm. Once, the algorithm learned to explore the topic it can recommend relevant websites based on its own explorations. These recommendations are then visited (and further explored) by the domain expert, again providing data for the algorithm to learn from. Over time the topic will be explored increasingly exhaustive.

It is also possible to use methods from the field of **transfer learning**, e. g. to 1) apply a learned model of one topic to another topic for the same expert, or 2) apply a learned model of one domain expert to another expert. With the help of exploration graphs we believe it is possible to develop more elaborate algorithms that are able to automate complex search tasks. The result of this learning goes beyond "learning to rank"-approaches described earlier, because it has to take into account the utility of information. That is, an automated search agent trained with data about search behavior of human domain experts should suggest only highly relevant and novel information.

7 CONCLUSION

This paper presents a hybrid log-based and observational approach for modeling search behavior. We formulate the problem of interpreting search behavior based on browser interaction logs and introduce the idea of an exploration graph to model transitional and semantic relationships during search. These interpretations can help to keep valuable information on how users explore an information space, if reasonable assumptions about the search behavior can be made. With the help of a user study we outline how semantic interpretations compare to interpretations without assumptions. Interpreting a user's interactions during search in an exploration graph may be key to various new investigations, e. g. how users interact in groups to fulfill a certain research task. Finding meaningful interpretations becomes a new challenge in the analysis of interaction logs. We also show how leveraging the data inherent in exploration graphs can be used in recommender systems to make search tasks in business settings more efficient. Since the quality of such recommendations depends heavily on the quality of the interpreted exploration graphs it is important to put a lot of effort into creating meaningful interpretations of search behavior in the first place. We believe that leveraging data about exploration graphs is a promising approach to tackle new research directions and produce highly innovative support systems, especially for professional searchers.

REFERENCES

Ageev, M., Guo, Q., Lagun, D., and Agichtein, E. (2011). Find it if you can: a game for modeling different types of web search success using interaction data. In *Proc. of the 34th Int. ACM SIGIR Conf. on R&D in IR*, pages 345–354. ACM.

Bates, M. J. (1989). The Design of Browsing and Berrypicking Techniques for the Online Search Interface. *Online Inform. Review*, 13(5):407–424.

Card, S. K., Pirolli, P., Van Der Wege, M., Morrison, J. B., Reeder, R. W., Schraedley, P. K., and Boshart, J. (2001). Information scent as a driver of web behavior graphs: Results of a protocol analysis method for web usability. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems*, CHI '01, pages 498–505.

Gwizdka, J. and Lopatovska, I. (2009). The role of subjective factors in the information search process. *J. Assoc. Inf. Sci. Technol.*, 60(12):2452–2464.

Hassan, A., Jones, R., and Klinkner, K. L. (2010). Beyond dcg: User behavior as a predictor of a successful search. In *Proc. of the 3rd ACM Int. Conf. on Web Search and Data Mining*, pages 221–230. ACM.

Knight, S. and Spink, A. (2008). Toward a web search information behavior model. In *Web Search*, pages 209–234. Springer.

Marchionini, G. (1997). *Information Seeking in Electronic Environments*. Number 9. Cambridge university press.

Marchionini, G. (2006). Exploratory search: From finding to understanding. *Commun. ACM*, 49(4):41–46.

Nürnberger, A., Stange, D., and Kotzyba, M. (2015). Professional collaborative information seeking: On traceability and creative sensemaking. In *Semantic Keyword-based Search on Structured Data Sources*, IKC 2015, pages 1–16, New York, NY, USA. Springer-Verlag New York, Inc.

White, R. W. and Drucker, S. M. (2007). Investigating behavioral variability in web search. In *Proc. of the 16th Int. Conf. on WWW*, pages 21–30.

White, R. W. and Roth, R. A. (2009). *Exploratory Search: Beyond the Query-Response Paradigm*. Synthesis Lect. on Information Concepts, Retrieval, and Services.

Wildemuth, B. M. and Freund, L. (2012). Assigning search tasks designed to elicit exploratory search behaviors. In *Proc. of the Symp. on HCI and IR*, pages 4:1–4:10.

Wu, W.-C., Kelly, D., and Huang, K. (2012). User evaluation of query quality. In *Proc. of the 35th Int. ACM SIGIR Conf. on R&D in IR*, pages 215–224. ACM.