The Impact of SearchTrails on the Quality of Collaborative Search

Sebastian Franken[1], Ulrich Norbisrath[2], Wolfgang Prinz[1]

Fraunhofer Institute for Applied Information Technology FIT, Sankt Augustin[1]
University of Applied Sciences Upper Austria, Hagenberg[2]

Summary

Several collaborative search systems build upon real-time collaboration during search processes. With SearchTrails, we present a novel way of capturing and exchanging the search process itself between collaborators. We achieve this by asynchronously exchanging search trails between collaborators and thus overcome the necessity of real-time interaction. In a study with 29 participants, we evaluate the value of search trails as collaboration artifacts to answer the research question whether search trails improve the quality of collaborative search results. We show that this is the case and users can build upon work of co-searchers in a very efficient way by analyzing and extending the given search trails.

1 Introduction

A number of approaches exist when it comes to the support of collaborative search processes. These approaches may be based on discussions (Morris & Horvitz 2007), preprocessed databases (Capra et al. 2012 & 2013), user-curated website collections (www.searchteam.com), or efforts towards supporting synchronous browsing processes (Golovchinsky 2008). These approaches build upon recommendations, tags, bookmarks, or user-defined sets of information. With SearchTrails, we realized an approach that goes beyond existing means of supporting collaborative search processes in that it captures the user's search process in its entity. Furthermore, our approach allows but not forces the user to enrich this information by valuable website excerpts (highlights) marking important, or by indicating explicitly negative search results. The search trails are stored on a remote server, such that users are able to recover the search results from other users and build upon them for their own search processes. We conducted a study with 29 participants for investigating the value of search trails for individual and collaborative search scenarios. In this paper, we address the research question whether the exchanging of search trails improves the quality of collaborative search results. We compare the exchanging of two types of artifacts with each other: On the one hand, written reports containing relevant information on a topic in prosaic form as a representative

Veröffentlicht durch die Gesellschaft für Informatik e.V. 2016 in
W. Prinz, J. Borchers, M. Ziefle (Hrsg.):
Mensch und Computer 2016 – Tagungsband, 4. - 7. September 2016, Aachen.
Copyright © 2016 bei den Autoren.
http://dx.doi.org/10.18420/muc2016-mci-0039

search result, and on the other hand the search trails containing a network of links and user-defined highlights that have led to the written reports. In our study, we compare building upon written reports as well as search trails for starting a search on a topic which is new for the participants. We analyze how the type of given artifact influences the number of visited resources and the quality of the artifacts resulting from the search process.

The work presented in this study builds upon the SearchTrails tool for supporting asynchronous, discontinuous, collaborative, and complex search processes. An early version of SearchTrails was evaluated in a qualitative study with users to show the effectiveness of our approach (Franken & Norbisrath 2014a) and was able to prove that search trails can be easily evaluated by evaluators (Franken & Norbisrath 2014b). A more refined version of SearchTrails allowed the exchange and the recreation of search trails and therefore the exchange of search processes between collaborators. With this version of SearchTrails, we could show by comparing usability metrics gained by the user experience questionnaire (UEQ) that search trail are superior to written reports when it comes to the exchange of search results (Franken et al. 2015). This version of SearchTrails was also used for the study described here.

In the following sections, we present underlying theoretical concepts and approaches for supporting collaborative search processes. We then describe the system overview and implementation of SearchTrails in more detail. An overview of the performed study and its results contribute the main part of this paper. The last section gives summarizes the findings and implications, and provides an outlook on future work with SearchTrails.

2 Theoretical background and related work

Complex search tasks, such as arranging a family holiday trip, colleagues researching for academic publications, or planning the construction of a technical device, involve a number of highly specific work steps, and are usually performed in several sessions, each building on previous knowledge. Complex search tasks as defined by Singer (2012) therefore consist of the search activities of aggregation, discovery, and synthesis. Furthermore, they tend to be multi-step and interactive processes, which are labor-intensive and time-consuming (Singer 2012). These properties refine the first definition of exploratory search processes by (Marchionini 2006), who classified search activities into ‘Lookup’, ‘Learn’, and ‘Investigate’, and constructed exploratory searches as the ones primarily involving the complex and cognitively demanding activities of ‘Learn’ and ‘Investigate’. To support complex search activities, tool support for the threefold of aggregation, discovery, and synthesis is necessary. However, current search engine support for these types of tasks is limited (Singer et al. 2013).

One approach for providing this support is the construction of search trails, which goes back to an idea by Vannevar Bush, who described search trails for the first time: ‘Thus he goes, building a trail of many items. Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item’ (Bush 1945). Bush furthermore describes exchanging search trails between collaborating searchers: ‘[He] photographs the whole trail out, and passes it to his friend for insertion in his own memex, there to be linked into the more general trail.' (Bush 1945). However, Bush's idea has not been turned into practice. Recent approaches try to make use of search trails, while they do not follow Bush's user-centric approach, but see the value of search trails more on server-side.

One approach based on a Microsoft Internet Explorer plug-in was followed by (Singla et al. 2010), where the authors distributed a plug-in collecting anonymous trails and sent them back for further analysis. The authors use a number of different algorithms for analyzing the trails with regards to length, breadth, depth, or diversity. While the authors claim that there is a 'value in trails', and hope that the best paths 'outperform the average over all trails followed by users', they do not perform a user evaluation where they return paths to searchers and evaluate their actual value. On a higher level, a study by (Awadallah et al. 2014) presents an approach where the query logs of the search engine 'Bing' are evaluated with regards to the IP addresses of users. The authors recreate the users' trails on the search engine and identify frequently visited clusters of pages. These search trails were used to generate recommendations for further user investigation.

Existing approaches for supporting collaborative search rely more on direct interaction of searchers, curated information collections, preprocessed databases, or ratings and recommendations. A collaborative search support system is presented by (Golovchinsky et al. 2008), where distinct roles of the searchers during the search process allow splitting work between collaborators. ResultsSpace (Capra et al. 2012 & 2013) relies on preprocessed databases and derives recommendations from the relevance ratings of its users. Recent approaches like content curation (Zhong et al. 2013) consider portals like Pinterest and Last.fm user-curated content collections, similar to the systems SearchTeam and Diigo. These systems allow users to generate and curate information collections on certain topics and to share them with friends or colleagues. However, these approaches do not capture the search process as a whole and therefore do not store the sidetracks of search processes, or the places where a user did not find relevant information. This so-called negative search (Garfield 1970) can be of great value, as it keeps collaborators from running into known dead ends.

Our approach builds upon a web browser extension, which enables users to actively capture their own search processes, including user-defined highlights and sidetracks. This has so far only been done for passive logging and evaluating search processes of single users. One example for this is the Wrapper framework (Jansen et al. 2006), which consists of a standalone application monitoring the user's interaction with the browser and the operating system. Wrapper was used to analyze the exploratory search behavior of users, as this was suspected to be a chaotic process, involving multiple systems and multiple episodes. Another example is the 'Search-Logger' system, which consists of a Firefox browser plug-in capturing search sessions of a single user (Singer 2012). The search sessions were analyzed by the authors to prove the existence of complex search as defined above. Unfortunately, none of these approaches tried to derive added value from these logs for the individual user or for enhancing collaboration during search processes.

Our system SearchTrails overcomes the limitations of existing related work, and combines key features from search logging and collaborative search support systems to provide support for asynchronous, complex search processes. SearchTrails therefore is the first system to investigate the individual value of search trails for collaboration support.

3 SearchTrails

SearchTrails aims for supporting asynchronous, discontinuous, collaborative, and complex search tasks by supporting aggregation, discovery, and synthesis and the exchange and recreation of search trails. SearchTrails supports aggregation by recording the search trail with all its visited web pages as nodes and side tracks, which helps keeping the context of search results. It supports synthesis by collecting valuable information pieces (highlights) in text-form from websites. These text blocks are collected in the highlights overview, and are stored with the node. Discovery is supported by visual representation of the search process as a force-directed graph as well as search term suggestions derived from the keywords of visited web pages (Fig. 1). These keyword suggestions guide searchers into related, but new search directions. The force directed graph visualization is a proven concept for visualizing large information collections (Eades & Huang 2000), and is used for the dynamic layout of search trails. SearchTrails regularly stores the user's search trail on a remote server to avoid losing data in case the extension is closed. This mechanism allows the recovery, recreation, and exchange of own and foreign search trails. The user-defined highlights are integrated in the search trail (blue nodes) and are displayed in a highlights overview, combining important information and its sources. Negative search results are indicated by purple nodes, and clusters (colored hulls) indicate nodes from the same host.

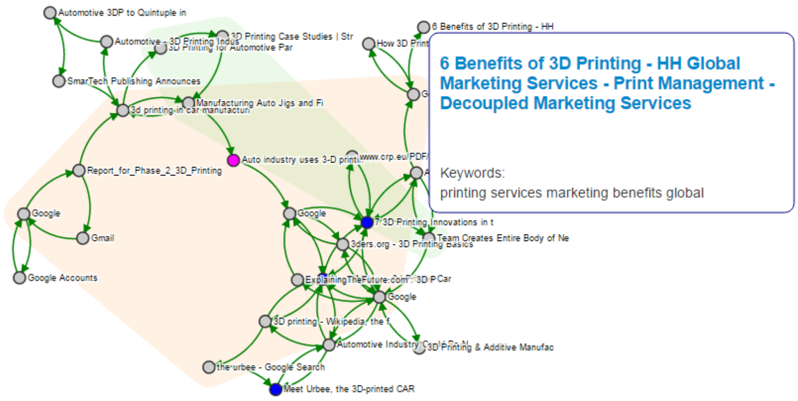

Figure 1: A user-generated search trail from the second study with users.

Literature and technical review suggest realizing SearchTrails as an extension for the Google Chrome browser, as it poses only minor installation effort and allows open-ended search processes, while unobtrusively logging the user's interactions with the browser. SearchTrails consists of three engines and the user interface on client side (Fig. 2), and the server-side infrastructure for saving and recovering the search trails. On client-side, the logging engine captures all events and transforms them into a JSON data structure, representing the search process. This data structure is stored in the browsers' background storage. The background storage reports about the changes of the data object, which triggers updates of the visualization by the visualization engine. The storage engine retrieves the search trail object from the background storage, and stores it with the help of server-side services. SearchTrails generates a unique ID for each search trail, which is stored on the user's hard drive and can be used for retrieving the search trail. This way, foreign users can retrieve the search trail data objects, given they know the ID. More details on the inner workings of SearchTrails can be found in (Franken et al. 2015).

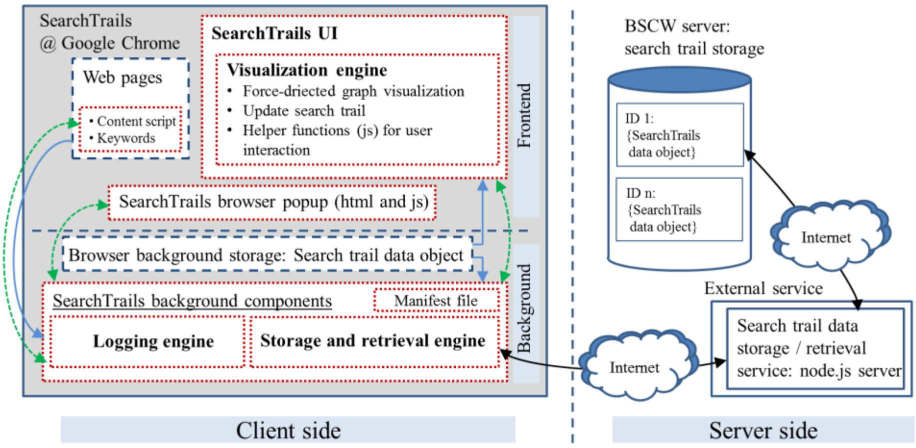

Figure 2: SearchTrails system architecture overview.

When users receive search trail IDs from colleagues, they can enter the ID into the retrieval mechanism. The server-side components retrieve the search trail data from the storage server and trigger its visualization. The visualization recreates the search trail, the keyword list, and the highlights without the nodes that were deleted by the previous user. For exploring an unknown search trail, users evaluate the search trail, the keyword list, and the highlight overview. All search trail nodes (see Fig. 1) can be dragged and rearranged in the force-directed layout. Hovering over a node reveals a window with the most important details on the visited URL. The search trail itself marks all nodes with added highlights in blue, while irrelevant nodes are marked in purple. The search trail is clustered by the visited hosts, and a cluster can be reduced to one larger cluster node, which replaces all nodes in the cluster. This way, users can evaluate search trails in a structured way, and detect more and less valuable clusters and add nodes to the search trail by visiting new web pages while SearchTrails is active.

4 SearchTrails evaluation study with users

For our study, we invited 29 students of a university lecture on Computer Supported Collaborative Work (CSCW) in a master course as a representative example of tech-savvy users with experience in web search. We provided detailed instructions on how to install SearchTrails and made sure the installation went well. For the study, we developed two search tasks to avoid biasing the study results by the selection of just one topic. The search tasks required the evaluation of given artifacts and checked to fulfill the seven characteristics for complex search tasks in (Kules & Capra 2009). The first search task covered 3D printing, while the second search task covered home automation. We found that the search tasks did not at all influence any of the search process results.

The study consisted of two phases of one week duration each, in which we divided the study participants in two groups, based on their technical support by SearchTrails and the type of given artifact to start from. During the first phase, group A used SearchTrails, while group B did not have SearchTrails available for searching. During the second phase, group A received a written report to start their search from, while group B received a search trail with similar information content as the written report. In each of the two groups, the participants worked on both topics to make sure that the results were not biased by any specific topic. The topics were handed out with alternating groups through the rows of the lecture hall. This prevented plagiarism and neighboring students being assigned to the same group. Between the two phases, the topics were exchanged between the participants, such that every participant started on an unknown topic. After the first phase, all participants were asked to produce a written report which contained all relevant information such that anyone who did not perform the search has an overview of the results. The participants from group A used SearchTrails and each produced a search trail with highlights on their topic. From these artifacts, we selected one average search trail and its corresponding report on each topic which contained a good basic set of information, leaving out more specific information about cars or home security. For the exchanged search trails, we made sure that they contained no personal information.

For the second phase, we supplied the participants with a search trail or a report from the first phase and asked them to build upon this material for answering the more specific questions. The participants who were given the search trail recreated it and started by evaluating it, while the participants with the reports started with reading them. This procedure resembles asynchronous collaboration on complex search tasks on a new and a traditional search process artefact. For 3D printing, we asked: ‘Based on the given material, find applications of 3D-Printing in the car manufacturing domain. Which applications exist, which ones will come? Will 3D-Printing change the way of manufacturing cars in the future?’. For home automation, we asked: ‘Based on the given material, find applications of home automation dealing with home security. Which applications exist, which ones will be available? Which applications would you prefer?’. We asked the participants to write a report on the new topic with all necessary information. Including URLs was not mandatory, and we did not request a minimum amount of text, to avoid the production of filler text. After one week, we collected the reports and stored the generated search trails in a safe place.

From our participants, we received 26 search trails and 21 reports. All three authors independently graded the participants’ reports with grades from 1 to 5 without knowing about the given artifact to objectively judge the quality of the reports (1 is the best grade and 5 means ‘failed’). As the reports were requested to be able to inform someone who did not perform the search about its key results, we graded the reports by their quantitative breadth and their qualitative depth of information. We similarly clustered the resulting search trails to resemble academic grades and performed statistical analyses on the generated search trails.

5 SearchTrails study results

We assume that when a collaboration artifact is given, the quality of the search result is more important than the quality of the search process. This is because the common goal of the collaborators is more on producing a collaborative result than in experiencing a highly qualitative search process. Therefore, our hypothesis deals with the quality of search results. In our case, the search results are the search trail and the report. From these search results, the search trail can be considered a direct result of the search process, while the report is an indirect result. First, we compare the average grades of the reports, depending on the given collaboration artifacts. We did not split the results by the given search topics, as analyses show that the topics themselves have no impact on the quality of the artifacts produced by each group. We performed statistical significance analyses on all discovered differences. While the average grade for group A, who received the report as collaboration artifact is 3.20, while the average grade for group B is 2.21. The results are highly significant on a 5% error level. They show that the average grade of the reports seems to heavily depend on the type of the given artifact. The reports of group B, whose participants were equipped with a search trail as a collaboration artifact, are graded approximately one full grade better than the reports of group A, whose participants were equipped with a report as collaboration artifact.

We also clustered the search trails from the second phase in an expert workshop. This clustering made use of the full spectrum of academic grades. The final clustering of the search trails shows that the search trails from group A spread around the full range of academic grades and achieve an average value of 3.17, while the search trails from group B achieve a statistically significant better average value of 1.78 (5% error level). This is not too astonishing when keeping in mind that the participants from group B were equipped with a proper search trail, which they should evaluate and enlarge, and the participants did not deteriorate the given search trail, e.g. by deleting nodes or highlights.

However, a statistical comparison of search trail characteristics reveals a number of significant differences between the two groups (Table 1). The values for group B show exclusively the value added during the search process for the second phase of the study; the values of the given search trails were already subtracted. The last two rows show whether the difference between the average values of group A and group B (added value) is significant on a 5% or a 10% error level. In many cases, the net value added to the given search trail by the participants from group B is significantly smaller than for group A.

The key properties of a search trail are the number of nodes, edges, steps, clusters, and highlights. Other key characteristics are the duration and the number of seconds not spent on search engine pages. The last characteristic is the average loop length. The numbers of nodes, edges, and highlights of a search graph are self-explaining. The number of steps is the number of user induced actions during the search process for navigation from one node to another. When a user walks the same path within a graph several times, the graph is not altered anymore, but the number of steps through the graph increases. The number of clusters is the number of hosts from which more than three different web pages were visited. The duration is the number of seconds in which the participants have searched actively, meaning that no interruptions of more than 15 minutes occurred. When two user-induced events are more than 15 minutes apart, the times are counted as idle times. We furthermore calculated where the participants have spent their time. The second to last column shows the number of seconds that were spent on non-search engine pages. The last column shows the average loop length. We count the length of all paths that start at a search engine until the path reaches a search engine page again and divide this by the total number of paths. This number serves as an indicator for the average depth with which the participants dived into the topic.

	# Nodes	# Edges	# Steps	# Clusters	# High-lights	Duration (s)	Time spent not on search engines (s)	Average loop length
Group A: Report	34.4	64.4	108.3	2.8	2.9	2752	1532	2.8
Group B: Search trail (added value)	23.0	28.4	42.8	1.0	1.0	1497	1006	5.0
T-test 5%	no	sig	sig	sig	no	no	no	sig
T-test 10%	no	sig	sig	sig	no	sig	no	sig

Table 1: Statistical data of search trails and their differences between the two groups.

Table 1 shows that all key characteristics of the search trails are on average larger for group A than they are for group B. These values show that the value that was added to the given search trail is significantly smaller for group B than for group A. This lower added value for group B did not occur randomly, but is significant in many cases: The lower numbers of edges, steps, and clusters are significant on a 5% error level, while the shorter duration is significant on a 10% error level. The participants from group B also produced significantly longer loops than the participants from group A. This means that these participants did longer tours through the Internet before going back to search engine pages and therefore dived deeper into the domain than the participants from group A. A search trail therefore seems to help avoid redundant searching and brings the users into a position where they are able to produce better results with lower efforts, meaning that they are significantly more efficient.

These results show that search trails as collaboration artifacts have an advantage over written reports as collaboration artifacts. They show that the participants who were given a search trail invested less resources into extending the given material than the participants who were given a report. These differences are statistically significant in most cases. Altogether, the collaborative efforts when extending the given search trail lead to significantly better search results. For group B, both the report and the search trail improve significantly. Adding to that, the results also show that the participants of group B were significantly more efficient: These participants invested only 54% of the time group A needed to invest into the search process, but ended up with a report that was on average one full grade better than from group A (3.20 vs. 2.21, see above). These results show that we can strongly confirm our research question: Exchanging search trails in collaborative search improves the quality of collaborative search results. SearchTrails eases collaboration and increases the efficiency of the searchers in collaborative search scenarios.

6 Conclusion and outlook

Our results from a field study with 29 participants over two weeks show that SearchTrails as a collaborative search support tool can induce significant improvements in terms of efficiency of collaborative work. This leads to a gain in the quality of collaborative search processes and validates our research question. Search trails as collaboration artifacts are proven to be a valuable means of exchanging search results and ease building upon previously done work.

In further analyses, we investigate the correlation between statistical search trail characteristics and the grades of the search trail and the report for all participants as one group. We can show that there are several characteristics that lead to an improvement of the search trail grade, which is not astonishing as these are the criteria we graded the search trails on. What is interesting is that even if we graded the report and the search trails independently, a number of search trail statistics show a positive correlation with the report grade. These characteristics are – among others – the number of visited non-search engine (NSE) pages, the number of steps through NSE pages, and the time spent on NSE pages.

This suggests that searchers with extensive search trails tend to produce good reports, and highly qualitative information from the search trail makes its way into the participants' reports, which indicates that a valuable search trail can be considered a head start into a topic. This may be based on the unfiltered insight into the collaborator's search process, its resources and valuable or less valuable results.

The results derived in this and the previous publications indicate that the concept of building, visualizing, exchanging, and evaluating search trails has an impact on collaboration for other systems. For the collaborating searchers, it seems to be easier to grasp the unfiltered contents of a collaborator's search trail than to evaluate a linear written text. This may be due to the playful nature of visual representations of complex search processes and their results, which can be evaluated individually. Furthermore, this contribution shows that it is more efficient to extend search trails and to write reports from this starting point than to directly extend written reports. The value lies in the exploration and the possibility to individually learn about the search trails contents. Our results also show a significant gain in efficiency when relying on the visual representation of a search trail compared to a written report. These results strengthen the value of visualizations of search processes and their results and should influence forthcoming systems.

Although the results of our study are positive, a potential flaw of our study is that the students are representative users of SearchTrails, but tend to do only necessary work, and are less intrinsically motivated. Especially that we did not set fixed guidelines for the reports or search trails to be delivered may not have improved the quality of the search results. However, as all students were affected by this motivational problem, the overall results in more real test cases will most likely be stronger than in our recorded samples.

Acknowledgements

We thank our study participants from the B-IT for the time they invested in our study and the reviewers for their helpful comments.

References

Awadallah, H., White, R.W., Pantel, P., Dumais, S.T. and Wang, Y.M. (2014). *Supporting complex search tasks*. Proc. of the 23rd ACM CIKM (pp. 829-838). ACM.

Bush, V. (1945). *As we may think*. The Atlantic Monthly 176 (July 1945). pp. 101–108.

Capra, R., Chen, A.T., Hawthorne, K. & Arguello, J. (2012). *ResultsSpace: An experimental collaborative search environment*. Proc. of the ASIS&T. 49 (1). pp. 1–4.

Capra, R., Chen, A.T., McArthur, E. & Davis, N. (2013). *Searcher actions and strategies in asynchronous collaborative search*. Proc. of the ASIS&T. 50 (1). pp. 1–10.

Eades, P. & Huang, M.L. (2000*). Navigating clustered graphs using force-directed methods*. J. Graph Algorithms Appl. 4 (3). pp. 157–181.

Franken, S. & Norbisrath, U. (2014a). *Supporting the evaluation of complex search tasks with the SearchTrails tool*. Proc. of 24th CASCON, Toronto, Canada: IBM Corp., pp. 262–274.

Franken, S. & Norbisrath, U. (2014b). *Trail Building During Complex Search Tasks*. In: Proc. of MuC 2014, München, Germany: De Gruyter Oldenbourg, pp. 135–144.

Franken, S., Norbisrath, U. & Prinz, W. (2015). *Search Trails as Collaboration Artifacts – Evaluating the UX*. In Proc. of MuC 2015, Berlin, Germany: De Gruyter Oldenbourg, pp. 23–32.

Garfield, E. (1970). *When is a negative search result positive?*. Ess. of an Inform. Scient. 1. pp. 117f.

Golovchinsky, G., Adcock, J., Pickens, J., Qvarfordt, P. & Back, M. (2008). *Cerchiamo: a collaborative exploratory search tool*. Proc. of CSCW 20008. pp. 8–12.

Jansen, B. J., et al. (2006). *Wrapper: An application for evaluating exploratory searching outside of the lab.* SIGIR Workshop on Evaluating Exploratory Search Systems.

Kules, B. & Capra, R. (2009). *Designing exploratory search tasks for user studies of information seeking support systems*. Proc. of JCDL '09, p. 419-420.

Marchionini, G.(2006). *Exploratory search: from finding to understanding*.Comm.ACM.49(4).p.41-46.

Morris, M.R. & Horvitz, E. (2007). *SearchTogether: an interface for collaborative web search*. Proc. of the 20th annual ACM symposium on User interface software and technology (pp. 3-12). ACM.

Singer, G. (2012). Web search engines and complex information needs (Dissertation). University Tartu.

Singer, G., Norbisrath, U. & Lewandowski, D. (2013). *Ordinary search engine users carrying out complex search tasks*. Journal of Information Science. 39 (3). pp. 346–358.

Singla, A., White, R. & Huang, J. (2010). *Studying trailfinding algorithms for enhanced web search.* In Proc. of the 33rd ACM SIGIR conf. on R&D in Information Retrieval (pp. 443-450). ACM.

Zhong, C., Shah, S., Sundaravadivelan, K. & Sastry, N. (2013). *Sharing the Loves: Understanding the How and Why of Online Content Curation*. In Int. Conf. on Web and Social Media.

Author information

Dr. Sebastian Franken & Prof. Wolfgang Prinz, PhD, Fraunhofer FIT, Sankt Augustin. {sebastian.franken, wolfgang.prinz}@fit.fraunhofer.de. Dr. Ulrich Norbisrath, University of Applied Sciences Upper Austria, Hagenberg. ulrich.norbisrath@fh-hagenberg.at.