Datenbank Spektrum
https://doi.org/10.1007/s13222-021-00375-x

DATENBANKGRUPPEN VORGESTELLT

The Search Studies Group at Hamburg University of Applied Sciences

Dirk Lewandowski[1] · Sebastian Sünkler[1] · Sebastian Schultheiß[1] · Helena Häußler[1] · Ulrike Spree[1] · Christiane Behnert[1]

Abstract
We present an overview of the work of the Search Studies research group, focusing on commercial search engines from a user perspective. This encompasses studying what users of these search engines get to see on the result pages, how users interact with search engines, and the effect both have on knowledge acquisition in society. Our research combines search engine data analysis, by collecting and analysing data from commercial search engines (data science), with understanding information-seeking behaviour through conducting user studies in different settings (information science), ranging from large, representative online surveys to behavioural studies in the lab employing, amongst others, eye-tracking.

Keywords Search engines · Search engine advertising · Search engine optimization · Information science research · Open Web Index · Information literacy · Data literacy · Search result analysis

1 Introduction

The Search Studies research group at Hamburg University of Applied Sciences (Hochschule für Angewandte Wissenschaften Hamburg) focuses on research on commercial search engines. We combine interests in studying what users of these search engines get to see on the result pages, how users interact with search engines, and the effect both have on knowledge acquisition in society.

Our group may be atypical for the information retrieval (IR) groups usually described in this section of *Datenbank Spektrum*, as we are not based in computer science but library and information science (LIS). LIS has a long tradition of investigating information-seeking behaviour in various settings and contexts [1–3]. We regard the LIS perspective as an essential addition to the computer science perspective, adding a deep understanding of user behaviour and a societal perspective.

The importance of understanding user behaviour, or the user perspective in more general, can be illustrated by quotes from two recent Salton Award winners, namely Kalervo Järvelin and Nick Belkin. Järvelin said in his acceptance speech that "A broader context is needed because optimal solutions may be missed if the R&D work on IR systems/services is not informed by their use context." [4, p. 56], and further, that the aim of IR research should be to "design technology that supports people as experts of their own work or life" (p. 61).

In his Salton Award acceptance speech, Nick Belkin stressed that developing user-oriented IR systems "requires the development of methods for investigation of the behaviors and intentions of people prior to, during, and after engagement in an IR system. Adapting methods of information behavior research to the specific requirements of IR system design is a likely path to such outcomes." [5, p. 2].

Building on that, understanding information (seeking) behaviour is at the heart of what we do: We first have to understand how users behave before designing interventions to help users become more information literate, or design systems that make it easier for users to find relevant information. Furthermore, we have to consider that search engine providers also know about how users behave and that they may try to subtly manipulate users to consider specific results [6, 7], not least the ones provided by the search engine company itself or advertisements shown in response to a query. So, focusing on commercial search engines (instead of working on experimental systems where researchers have full access to the database and the algo-

✉ Dirk Lewandowski
dirk.lewandowski@haw-hamburg.de

[1] Faculty Design, Media, Information; Department of Information, Hochschule für Angewandte Wissenschaften Hamburg, Finkenau 35, 22081 Hamburg, Germany

Published online: 22 June 2021

rithms) requires specific methods and approaches. We address this by combining

- *Search engine data analysis*, collecting and analysing data from commercial search engines (data science)
- *Understanding information-seeking behaviour* through conducting user studies in different settings (information science), ranging from large, representative online surveys to behavioural studies in the lab employing, amongst others, eye-tracking.

We firmly believe that only by combining these two will we fully understand what search engines offer to their users and how users interact with the systems and the information presented to them. Such a research program, of course, requires methods from different disciplines, namely computer science and the social sciences. It is interdisciplinary and requires researchers not only to be able to work together with researchers from other disciplines but for themselves to have an interdisciplinary profile. Information science has traditionally used a variety of methods, mainly from the social sciences, and has increasingly been using computer science methods.

In the remainder of this paper, we provide an overview of the range of methods applied in our research, and then present some of our current research projects, some of our collaborations with industry and other institutions, and our involvement in teaching at the Department of Information.

2 Methods

The approach described in the introduction, of course, requires the use of multiple methods, deciding which method(s) to use for every research question anew. We deem it important to introduce methods here because the strengths of these methods (and their combination in multi-methods studies) may not be known in depth in the different communities of researchers we address. Of course, we can only give a brief overview here:

- Eye-tracking: Eye-tracking in a lab setting allows us to collect data on users' visual attention on search engine result pages (SERPs; for an overview, see [8]). We can see to which results users are attracted and what they leave out when examining the SERPs. Information on both is crucial as users will only select results that they have attended to before.
- Large-scale, representative online surveys [9, 10]: For descriptive studies, valid samples of the targeted population are essential. As commercial search engines are used by most of the people using the Internet [11], we need to aim for samples representing the Internet population. This is not feasible without working together with market research firms with access to large panels of Internet users. However, we found that conducting this type of studies is not only feasible but also manageable in terms of research budget. The outcome of such large studies is, of course, of high evidential value and complements other types of studies.
- Experiments: Experiments (as defined in the social sciences) are explanatory studies that allow for a causal conclusion to be drawn. However, experiments have to be designed carefully to obtain valid results. Given such careful design, they can provide high evidential value and, in that way, improve and strengthen theory.
- Transaction-log analysis: To understand how people behave when using search systems, we analyse transaction logs. It is not easy to gain access to such logs, but our long-lasting collaboration with different providers of search systems makes this a lot easier.
- Screen-scraping and automatic data analysis: When investigating search engines, it is paramount to understand what results users see in their searches. To this end, we automatically query commercial search engines, screen-scrape their result pages, and analyse the results found.
- Classification and machine learning: For many research studies, search results have to be classified according to specific criteria. We use manual classification of large numbers of results to test our systems and train machine learning algorithms.
- In addition to experimental and quantitative methods, we also apply qualitative methods observing users interacting with a system in real-life situations, e.g., in contextual inquiry, and interview different stakeholders like the users themselves and other actors like search engine providers or SEO specialists.

However, while each method described can be used to address specific types of research questions only, the vital point is to combine methods in a meaningful way and thereby not only add one to the other but create something new (through multi-methods study designs). In the following sections reporting on our results, we will show that a combination of research methods often allows for a deeper understanding of a topic.

3 Influences on Search Results in Commercial Search Engines

One of our main research interests is to understand user interactions in commercial search engines. A pivotal point is first to know how search engine result pages are assembled from different sources and what types of results users are confronted with in their searches [12]. In Fig. 1, different actors influencing what is seen on the SERPs are

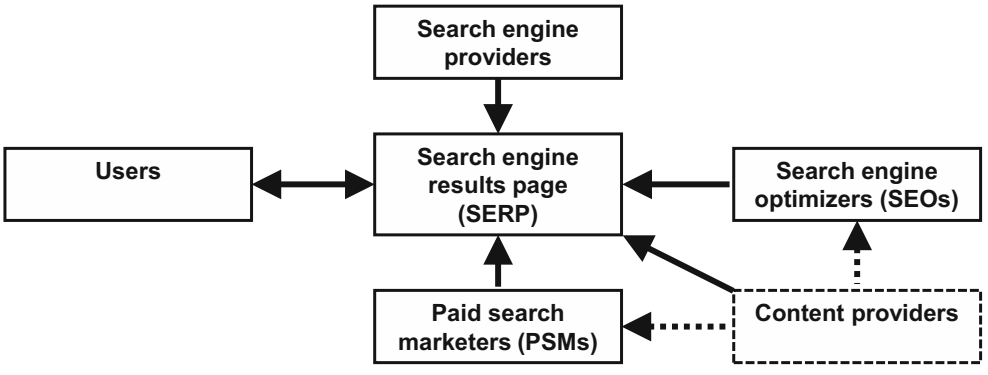

Fig. 1 Influences on the search engine result pages (from [10])

shown. While the search engine provider obviously has the most extensive influence, also users (through their clicks and reading behaviour), as well as content providers, influence what is shown. As visibility in commercial search engines, especially Google, is of high importance to these content providers, they use the services of paid search marketing (PSM) to place ads on the SERPs and of search engine optimizers (SEOs) to make their content more visible in the organic results section of the SERPs. In the following, we will briefly review our research in these areas.

3.1 Understanding the Role and Effect of Search-based Advertising

A basic assumption of information retrieval research is that the aim of developing IR systems is to provide users with the best possible results [13–15]. However, in a commercial setting where users do not pay for using the IR system, but system providers make money from contextual advertising ("sponsored results"), these providers may be tempted to blur the lines between ads and organic results to increase their ad revenue.

Ads in a search engine are particular in that they are shown in response to user queries. Therefore, these ads may be relevant to the query and can be regarded as a result type [9]. However, the problem with this form of advertisements is that many users cannot necessarily distinguish between these and the so-called organic results, i.e., results that are shown without the search engine company being paid for.

We address the question of how contextual advertisements influence information seeking by a combination of methods: We conducted a survey using a representative sample of the German online population in which we asked questions about users' understanding of advertisements in Google (e.g., whether they know how Google generates the most significant part of its revenues, how one can distinguish ads from organic results) and combined that with tasks where participants had to mark ads or organic results on SERPs they were shown [9]. Furthermore, in an experimental setting [16], we divided participants into two groups. Both groups were shown the same search results, but in one group, the top two results were labelled as ads, whereas in the other group, there were only organic results.

We found that most users are not knowledgeable of Google's business model and cannot identify ads on SERPs. In the experiments, those users not knowledgeable of ads click the first-shown ad more than twice as often as other users. This means that these ad clicks are not informed decisions, but users select them expecting to click on an organic result.

In an eye-tracking study [17], we confirmed the study's results and added evidence that users with little knowledge of ads are less willing to scroll down to the organic results and that ads receive significantly more visual attention when displayed on small screens (e.g., smartphones) than on large ones (e.g., desktop PCs).

The consequences for IR research are that the commercial interests of IR system providers should be taken into account and that, especially in interactive information retrieval (IIR), research studies should consider that ads may influence user search processes. Studies where ads are not considered risk producing results that only hold in the lab but lack external validity.

It should be stressed that contextual advertising in search systems does not only occur in general-purpose search engines like Google but also in other types of search systems like e-Commerce sites such as Amazon or eBay. On both these platforms, the advertisements shown on result pages are labelled as "sponsored".

3.2 Understanding and Measuring the Effect of Search Engine Optimization (SEO)

Apart from paid search marketing, external actors can influence the results of commercial search engines by trying to gain visibility in the organic results through search engine optimisation measures (SEO). Search engine optimization is "the practice of optimizing web pages in a way that improves their ranking in the organic search results" [18]. While SEO is usually associated with optimizing websites

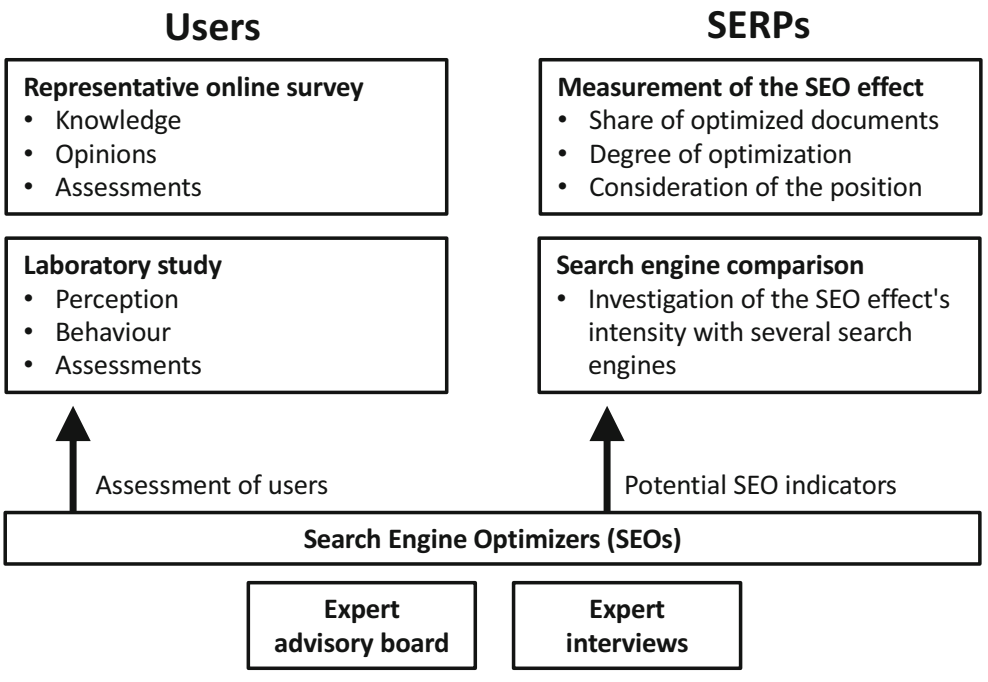

Fig. 2 SEO-Effekt research framework (from [19])

to promote products and services, it reaches far beyond that: SEO is also applied to information sources (websites) relating to topics of societal relevance, e.g., politics and health.

We started to address the topic of SEO in the research project "SEO-Effekt" funded by the Deutsche Forschungsgemeinschaft (German Research Foundation, DFG). The project aims to gain a fundamental understanding of SEO from the users' point of view and to build software that reliably detects SEO efforts on web pages, therefore making it possible to show how the (top) results of commercial search engines are influenced by SEO, and how certain types of results (e.g., non-commercial, documents from government websites) may get demoted in the result lists.

As can be seen from Fig. 2, the project consists of two parts, users and SERPs. Both are informed by expertise from the project's expert advisory board and interviews with experts from SEO, content providers and journalists familiar with SEO techniques [20]. In the user part of the project, we conducted a survey representative of the German online population, similar to the survey reported in the section on ads above. This survey again consisted of questions on users' knowledge, this time on advertising practices and SEO. In the task-based part, participants were asked to identify areas on the SERP that could either be influenced through SEO or paid search marketing. The overall result from the survey is that users do not have much knowledge about SEO (even less than about PSM). This result is important as it shows that, while there are huge investments in SEO (see [21]) to influence the organic results of commercial search engines, users understand little how this might affect the results they get so see. Furthermore, in an analysis of trust in search engines in general and Google in particular, we found that users with little search engine knowledge are more likely to trust and use Google than users with greater knowledge [10].

In the data analysis part of the project, we built software to query commercial search engines and extract information from their result pages. We further implemented automatic analysis to detect indicators for a page which is optimized for search engines. These indicators were derived from the professional SEO literature (e.g., [22, 23]) and from the expert interviews conducted within the project [20].

Our model currently consists of 48 factors, which can be grouped along three dimensions: tools and plugins, URL lists, and indicators for SEO. We prioritized these factors for the implementation in our system. Our approach combines the automatic identification of SEO indicators from web pages and websites with manually generating lists of optimized and not optimized websites. We follow this approach as results in search engines are not equally distributed, i.e., there is only a relatively small set of websites that account for a large fraction of the URLs shown [24] and clicked in the top results [25]. This means that, by manually classifying a limited set of websites, we can already detect a relatively large fraction of optimized pages. We created five lists of manually classified websites: SEO customers, news websites, online shops, business websites, websites with ads, and not optimized websites. In the automatic analysis, we focused on identifying tools used by search engine optimizers on the one hand and indicators for SEO on the other

hand. We differentiate between two types of tools: Tools particularly used for search engine optimization and analytics tools that are not necessarily used for SEO purposes exclusively but are usually used in the SEO context. We identified tools through analysing the HTML code of the results found. When a tool is used, a hint can generally be found in the HTML comments or a script. In terms of SEO indicators, we extract information from the page's HTML code, further information on the website level, and test for page speed. Data extracted from the HTML code includes the use of a page description and *nofollow* links, among others. Information on the website level includes the use of SEO-specific information in the robots.txt file and the use of a sitemap file. Finally, we measure page speed, as one of the fundamental technical factors in search engine optimization is optimizing the pages to load quickly. We used these indicators to build a rule-based classifier to determine the probability of search engine optimization in four classes: definitely optimized, probably optimized, probably not optimized, and definitely not optimized.

First results based on a dataset of more than 250,000 results show that the vast majority of results found in Google are optimized. However, it should be noted that our analysis is restricted to the limited number of results Google gives back in response to a query. While the estimated number of results Google shows on top of the result list often indicates millions of results, in reality, if one clicks from result page to result page, only a few hundred results are shown. Nevertheless, considering these results, it is striking that 88 per cent of results are classified as at least probably optimized. Furthermore, when considering the result positions, it is interesting to see that, while the probability of a result being optimized decreases with the result position, this probability is still about 80 per cent at position 100 [26].

These results show that users are confronted with a vast proportion of results that have been optimized for visibility in search engines. In future research, we will investigate how these optimized results might demote high-quality information from pages that have not been optimized, e.g., government websites. In future research, we will also improve indicator detecting. This will be achieved through the use of machine learning. Initial tests with common classification algorithms such as Decision Trees, Naive Bayes, and Support Vector Machines are promising. Our goal is to get a more accurate estimate on the SEO probability and to reduce the relevance of indicators that we have to manually maintain, such as the SEO tool lists and the classification of web pages in the overall evaluation, as we see a certain error-proneness there [26].

SEO is most often associated with general-purpose web search engines. Still, it also is applied to, for instance, academic search engines [27], with even publishers giving authors tips on how to improve their papers in terms of SEO. For IR research, it is essential that in IR systems where the database is not in complete control of the system provider, external actors will probably try to use them to gain visibility. Therefore, improving a system might not only include optimizing relevance measures but also detecting attempts to game the ranking algorithms.

3.3 Understanding the Role and Effect of Search Engine Providers' Self-interests

It is evident that the owner of a search system has the largest influence over its results. What is new in commercial search engines is that their providers usually not only offer search functionality but some other services that will profit from visibility in that search engine. That means that apart from applying SEO techniques to their own products, search engine companies can also prefer their own offerings in their search results. This can apply to organic results where, for instance, one can ask why results from YouTube are so prominent in Google's search results. Of course, YouTube is the world's most popular video platform. Still, it is also a subsidiary of Google. One could ask which part of its popularity has been established by showing its results in Google's top positions and its video results.

Apart from organic results, it has been found that Google prefers results from its shopping search engine to those from competitors. Results from shopping search are integrated into SERPs as vertical results, i.e., boxes with results from special collections. A competitive investigation by the European Commission led to the Commission fining Google the then record-breaking sum of 2.4 billion Euros for preferring its shopping results over that of competitors. In the context of this competitive investigation, Google provided the Commission with several proposals for displaying their shopping results along with results from competitors. We conducted several studies investigating the effect of these proposed changes, using representative samples of 1000 participants per country investigated [28]. We found that the proposed changes did not lead to a shift in user behaviour, satisfying the need for fair representation of competitors.

While the studies reported are commissioned reports, they have been independently conducted. Not lest for basing regulatory decisions on sound evidence, it is essential to conduct such large-scale, representative studies of the relevant online population. Companies like Google, while of course having such data, do not provide the public with data on searcher behaviour.

4 Understanding User Behaviour, User Queries and Improving Rankings in Library Search Systems and the Academic Context

As our background is in library and information science, a recurring interest has been in library search systems. In the LibRank project, a collaborative project with the ZBW Leibniz Information Centre for Economics, funded by the Deutsche Forschungsgemeinschaft, we investigated how new approaches to result ranking can be applied in library search systems. Here, we primarily focused on incorporating popularity factors like the number of citations and the number of downloads. Our large-scale evaluation of the different ranking approaches found that there is no single-best way to rank results in library search systems and that popularity factors do not improve the rankings considerably from a user point of view. Another outcome of the project is a demonstrator where one can adjust the ranking algorithm based on the relevance factors used in the project and re-rank results from the EconBiz library search system accordingly (see https://github.com/LibRank-Project).

In the context of searches in library search systems, we also conducted several transaction log analyses based on more than 4 million search sessions from the library search systems EconBiz and Beluga [29]. Of particular interest here was how these systems handle known-item searches (see also [30]) and what search tactics users apply. As reasons for unsuccessful known-item searches, we identified incorrect or ambiguous queries, erroneous metadata, and an item not being in stock. Moreover, we found that users employ a variety of search tactics, such as changing the query length, with varying success in their searches. In addition, probably due to the lack of spell-checking functions in the tested systems, frequent corrections of erroneous queries were noted.

While library search systems can be considered the "original information retrieval systems" [31, 32], much has to be done to understand the specifics of searching these systems and ranking their specific type of material, even more so as they often combine short representations of physical material with full-text material. Furthermore, as with most other information retrieval systems, they cannot be regarded as standalone systems. Instead, they are one system integrated into a wider information environment where users move back and forth between general-purpose web search engines and more specialized search systems [33].

5 Other Information-Seeking Behaviour Studies

Regarding the broader context of academic search systems, it is crucial to understand what criteria users apply to decide whether a search result is helpful to solve their search task or not, while acknowledging the fact that such criteria are diverse and highly subjective. The purpose of the project "Relevance Clues" is to determine subjective criteria and the interdependence between different criteria on individual relevance judgments regarding search results in an academic context, i.e., by interacting with academic search systems such as Google Scholar or the ACM Digital Library. Dependable results will be achieved by applying an experimental research design involving human test persons—an approach still rarely used in information science—and by building a sample large enough for statistically firm statements. Results so far include a user model for relevance assessments in academic search systems taking into account popularity data, e.g., citation or download frequencies of the particular work [34]. The model aims at a systematic representation of elements of a search result as potential relevance clues, subjective relevance criteria users apply when assessing the results, and contextual relevance factors as variables influencing relevance judgments. Furthermore, a methodological framework for incorporating popularity data in experimental studies on relevance criteria has been developed [35].

As we see experimental designs—understood, as in psychology, as a method for gaining causal explanations—as a fruitful way to improve our understanding of IR systems, we aim at a more profound insight on what influences users selecting results in search engines.

In terms of users attempting to recognise credible information among the vast amount of search results, it is important to get an understanding of the role of search engines themselves and the long lasting and close relations users maintain with them. The presence of a reputed and trusted search engine brand may activate certain heuristics and help to reduce uncertainty and cognitive effort during the search and selection process. The aim of the research project "Selection decisions under uncertainty" is to determine whether there is a credibility transfer from commercial search engines to the displayed search results—an issue that has not been addressed so far. The research question should be addressed with an experimental research design on the one hand and focus group discussions on the other hand. Experimental conditions will include search engine brand, ranking and credibility of search results. The expected results are supposed to inform a model of the factors influencing the search process and hints for including relation aspects into information literacy education.

In terms of approaches trying to integrate system- and user-based approaches to information retrieval evaluation, we developed a framework for combining session-based logging (e.g., [36]) with user relevance assessments after the session [37]. Thereby, we can have users assess the results they encountered during the session but also inject results, e.g., from search engines they did not use in their session, to compare search engines based on real user queries instead of using queries pre-defined by researchers.

6 Data and Information Literacy

Our and many other empirical studies have shown, users tend to overestimate their information literacy regarding the use of search systems. As we are interested in the effect that commercial search engines have on knowledge acquisition and findings show that users' competence in understanding and using these systems is relatively low, the question arises on how to improve this situation. Therefore, within the research group we broaden our scope on aspects on different literacies when dealing with information and data in general. We need to understand people's information behaviour and how they use and make sense of data and data products. Information literacy encompasses of course much more than the competent use of search engines. The discussion on the different literacies (e.g., media literacy, digital literacy, statistical literacy) sheds light on the different communities of practice involved in fostering an adequate dealing with data, information and knowledge and opens different access points to support learners to acquire the needed skills and competencies. Whereas, for example, the library and information science community focuses on competencies concerned with managing data, especially research data in all stages of the data lifecycle, the business community stresses the aspect of supporting decision based on data. Independent of the special need in data and information competencies within different disciplines and communities of practice, we adopt an integrated concept that comprises basic skills to deal with data and information in a competent way, and for integrating such literacy education on a generic level into university education. Our activities consider curricula as well as extra-curricular interventions. The Erasmus+-funded international project "Data Literacy in Context" (DaLiCo) brings together four European Universities of Applied Sciences with the mutual mission to foster data literacy education: HAW Hamburg as the Coordinator of the project, Stichting Hogeschool Utrecht, University of Debrecen and Universitat Politècnica de València as partners. DaLiCo is focusing on increasing the visibility, quality and usage of existing data literacy activities at the participating universities. The objectives of the project result in five interconnected sub-projects: the compilation of data literacy resources leading to a so-called Data Literacy Map, a tried and tested modular Train-the-Trainer concept to stimulate the enhancement of data literacy competencies and allow exchange, a concept for data literacy learning spaces, a measurement and assessment tool for data literacy activities, and finally a pilot concept for the regional implementation of common international data literacy standards of the partners.

7 Real-World Impact

Our research has many practical implications, and research results can be applied outside academia. In our view, research should have an impact on the real world. Therefore, we cooperate with industry as well as institutions. Regarding industry, we mainly cooperate with companies through teaching projects and B.A. and M. A. theses (see "teaching" section below). However, in some industry-funded projects,

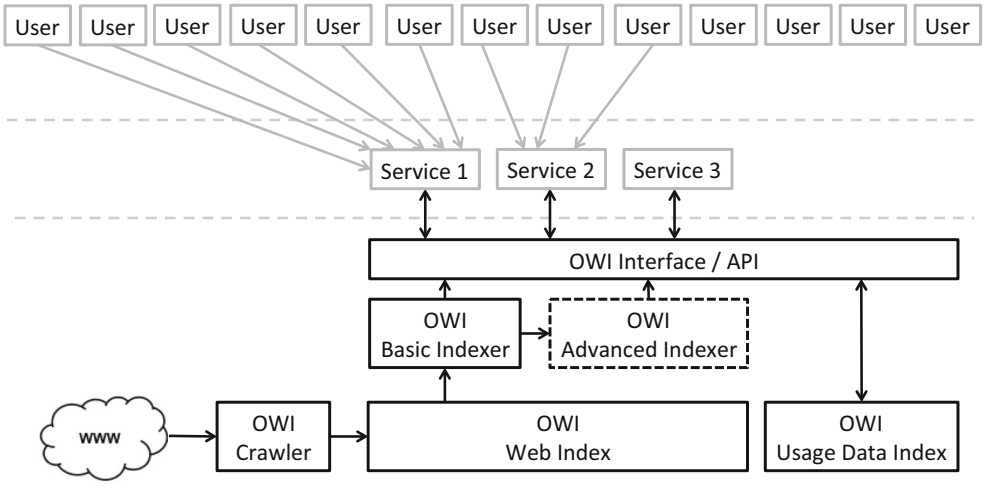

Fig. 3 Open Web Index (OWI) infrastructure

we also conducted studies at the intersection between search and user experience, an area that is not well understood.

An obvious implication of our research on commercial search engines is that some form of regulation of search engines is needed, especially when there is one such engine dominating the market. It should again be stressed that search engines are crucial for knowledge acquisition in society. According to the European Union, they are to be considered critical infrastructures (for a discussion, see [2]). In this regard, we are in continuous contact with media regulation authorities and government bodies. Still, as reported above, we also conducted commissioned studies for a plaintiff in the context of the European Commission's competitive investigation against Google.

Other outreach activities include giving expertise to the media and activities promoting our research results and, more generally, the need to understand search engines better.

Last but not least, we proposed the idea of establishing an Open Web Index (OWI) [38–40] as public infrastructure. The basic technical idea is to separate the index from the services that use the index, thereby allowing multiple search engines and other services to be built on the index (see Fig. 3). This index would be public infrastructure, while the services could be proprietary.

8 Teaching

The Search Studies group is part of the Department of Information at HAW Hamburg. Members of the group are involved in teaching in two study programmes: Library and Information Management (B.A.) and Digital Transformation (M. A.). In the reformed B.A. programme starting in the fall of 2021, students can choose from three profiles, one of which is Search and User Experience. Our expertise in the fields of user experience, knowledge representation and various aspects of search allows for a holistic perspective on search systems and their application and use.

Apart from required courses on user experience, knowledge organization, information retrieval, search engine optimization and information research, we provide elective courses on these topics. Specific to universities of applied sciences is their orientation towards practice and their focus on project-based learning. These specifics have led to many fruitful collaborations where students work on specific problems from industry. For instance, we taught courses where students worked on improving search and usability of the search systems of T-Online and Otto. Our project-based teaching has already won multiple awards, honouring our approach of combining quality teaching with close cooperation with industry and institutions. Awards received include the Ditze Award and the TIP Award ("Team Award Information Professional").

Study programmes at our department do not primarily address a technical audience, and our topics and approach differ from the many information retrieval courses taught in computer science. We address students who will need to better search, manage IR systems, focus on the user side of information retrieval, and understand the role of search (engines) for knowledge acquisition in society. The courses we offer are, therefore, less technical and have a broader focus, also towards a more social-science approach, than the curricula taught or proposed within the IR community (e.g., [41] and [42]).

As German universities of applied sciences are not allowed to award PhD's, we established cooperation with two research universities, namely Humboldt-Universität zu Berlin and the University of Hildesheim, to jointly supervise PhD students.

9 Open Science Practices

As a reaction to the replication crises in psychology but also more generally as a way to foster "science done right", the Open Science movement has advocated for not only making research papers freely accessible without cost to readers (Open Access) but also for making all research data, software code and the documentation of research procedures available. We agree that open science practices will help improve the replicability and evidential value of scientific results and, therefore, follow these practices in our research projects. While we have been making research papers open access for years, in 2019, we also started to make all research data and software code openly available through Open Science Framework repositories and other platforms. Further, as research procedures and software development often cannot be described in full detail in formal publications, we decided to apply a working-paper-based approach. Within a project, we document everything in working papers and make them available to the research community.

10 Conclusion

This article described the approach our research group is taking on researching commercial search engines from a user perspective. By describing some of our research projects, we showed how researchers could produce meaningful research even when they do not have full access to the search system (see [43]), as is usually the case in information retrieval research. Our approach to combining computer science, information science, and social sciences allows us to understand users' search engine use and its

impact on their knowledge acquisition through these engines. We hope that more researchers will see investigating commercial search engines as fruitful and worthwhile. We look forward to working with the research community on this.

Further information on our work, including links to data, publications, software and demos, can be found at https://searchstudies.org.

Acknowledgements We would especially like to thank our student assistants, past and present, whose work often remains unrecognized in research outputs.

References

1. Buckland M (2017) Information and society. MIT Press, Cambridge
2. Haider J, Sundin O (2019) Invisible search and Online search engines. Routledge, Oxford, New York
3. White RW (2016) Interactions with search systems. Cambridge University Press, New York
4. Järvelin K (2019) Salton award keynote. ACM SIGIR Forum 52:52–63. https://doi.org/10.1145/3308774.3308782
5. Belkin NJ (2015) Salton award lecture people, interacting with information. In: SIGIR 2015—Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM Press, New York, pp 1–2
6. Lewandowski D (2014) Die Macht der Suchmaschinen und ihr Einfluss auf unsere Entscheidungen. Inf Wiss Prax 65:231–238. https://doi.org/10.1515/iwp-2014-0050
7. Lewandowski D, Kerkmann F, Sünkler S (2014) Wie Nutzer im Suchprozess gelenkt werden: Zwischen technischer Unterstützung und interessengeleiteter Darstellung. In: Stark B, Dörr D, Aufenanger S (eds) Die Googleisierung der Informationssuche – Suchmaschinen im Spannungsfeld zwischen Nutzung und Regulierung. De Gruyter, Berlin
8. Lewandowski D, Kammerer Y (2020) Factors influencing viewing behaviour on search engine results pages: a review of eye-tracking research. Behav Inf Technol. https://doi.org/10.1080/0144929X.2020.1761450
9. Lewandowski D, Kerkmann F, Rümmele S, Sünkler S (2018) An empirical investigation on search engine ad disclosure. J Assoc Inf Sci Technol 69:420–437. https://doi.org/10.1002/asi.23963
10. Schultheiß S, Lewandowski D (2021) Misplaced trust? The relationship between trust, ability to identify commercially influenced results, and search engine preference. J Inf Sci. https://doi.org/10.1177/01655515211014157
11. European Commission (2016) Special eurobarometer 447—Online platforms. European Commission, Brussels https://doi.org/10.2759/937517
12. Lewandowski D (2018) Suchmaschinen verstehen, 2nd edn. Springer, Berlin, Heidelberg, New York
13. Manning CD, Raghavan P, Schütze H (2008) Introduction to information retrieval. Cambridge University Press, Cambridge
14. Croft WB, Metzler D, Strohman T (2009) Search engines: information retrieval in practice. Pearson, Boston
15. Büttcher S, Clarke CLA (2010) Information retrieval: implementing and evaluating search engines. MIT Press, Cambridge
16. Lewandowski D (2017) Users' understanding of search engine advertisements. J Inf Sci Theory Pract 5:6–25. https://doi.org/10.1633/JISTaP.2017.5.4.1
17. Schultheiß S, Lewandowski D (2021) How users' knowledge of advertisements influences their viewing and selection behavior in search engines. J Assoc Inf Sci Technol. https://doi.org/10.1002/asi.24410
18. Li K, Lin M, Lin Z, Xing B (2014) Running and chasing—The competition between paid search marketing and search engine optimization. In: Proc Annu Hawaii Int Conf Syst Sci, pp 3110–3119 https://doi.org/10.1109/HICSS.2014.640
19. Schultheiß S, Lewandowski D (2021) Expert interviews with stakeholder groups in the context of commercial search engines within the SEO effect project. https://osf.io/5aufr/. Accessed: 9 June 2021
20. Schultheiß S, Lewandowski D (2020) "Outside the industry, nobody knows what we do" SEO as seen by search engine optimizers and content providers. J Doc 77:542–557. https://doi.org/10.1108/JD-07-2020-0127
21. McCue T (2018) SEO industry approaching $80 billion but all you want is more web traffic. In: forbes.com. https://www.forbes.com/sites/tjmccue/2018/07/30/seo-industry-approaching-80-billion-but-all-you-want-is-more-web-traffic/. Accessed: 9 June 2021
22. Erlhofer S (2020) Suchmaschinen-Optimierung: das umfassende Handbuch, 10th edn. Rheinwerk Computing, Bonn
23. Enge E, Spencer S, Stricchiola J (2015) The Art of SEO: mastering search engine optimization, 3rd edn. O'Reilly, Sebastopol
24. Petrescu P (2014) Google organic click-through rates in 2014. https://moz.com/blog/google-organic-click-through-rates-in-2014. Accessed: 9 June 2021
25. Goel S, Broder A, Gabrilovich E, Pang B (2010) Anatomy of the long tail. In: Davison BD, Suel T, Craswell N, Liu B (eds) Proceedings of the third ACM international conference on Web search and data mining—WSDM '10. ACM Press, New York, p 201
26. Lewandowski D, Sünkler S, Yagci N (2021) The influence of search engine optimization on Google's results: a multi-dimensional approach for detecting SEO. In: WebSci '21: Proceedings of the 13th ACM Conference on Web Science. ACM, New York https://doi.org/10.1145/3447535.3462479
27. Beel J, Gipp B, Wilde E (2010) Academic search engine optimization (ASEO/span). J Sch Publ 41:176–190. https://doi.org/10.3138/jsp.41.2.176
28. Lewandowski D, Sünkler S (2013) Representative online study to evaluate the commitments proposed by Google as part of EU competition investigation AT. 39740-Google: report for Germany. https://searchstudies.org/wp-content/uploads/2021/06/Google_Online_Survey_DE.pdf. Accessed: 9 June 2021
29. Schultheiß S, Linhart A, Behnert C et al (2020) Known-item searches and search tactics in library search systems: results from four transaction log analysis studies. J Acad Librariansh 46:102202. https://doi.org/10.1016/j.acalib.2020.102202
30. Behnert C, Lewandowski D (2017) Known-item searches resulting in zero hits: considerations for discovery systems. J Acad Librariansh 43:128–134. https://doi.org/10.1016/j.acalib.2016.12.002
31. Sanderson M, Croft WB (2012) The history of information retrieval research. Proc IEEE 100:1444–1451. https://doi.org/10.1109/JPROC.2012.2189916
32. Stock WG, Stock M (2013) Handbook of information science. De Gruyter, Berlin
33. Lewandowski D (2018) Zugänglichkeit von Information Services und ihren Inhalten über Suchmaschinen. In: Schade F, Georgy U (eds) Praxishandbuch Informationsmarketing. De Gruyter, Berlin, Boston, pp 358–369
34. Behnert C (2019) Investigating the effects of popularity data on predictive relevance judgments in academic search systems. In: Proceedings of the 2019 Conference on Human Information Interaction and Retrieval. ACM, New York, pp 437–440
35. Behnert C (2019) Kriterien und Einflussfaktoren bei der Relevanzbewertung von Surrogaten in akademischen Informationssystemen. Inf Wiss Prax 70:24–32

36. Singer G, Norbisrath U, Vainikko E et al (2011) Search-logger analyzing exploratory search tasks. In: Proceedings of the 2011 ACM Symposium on Applied Computing—SAC '11. ACM Press, New York, pp 751–756
37. Sünkler S, Lewandowski D (2017) Does it matter which search engine is used? A user study using post-task relevance judgments. Proc Assoc Inf Sci Technol 54:405–414. https://doi.org/10.1002/pra2.2017.14505401044
38. Lewandowski D (2019) The web is missing an essential part of infrastructure: an open web index. Commun ACM 62:24–27. https://doi.org/10.1145/3312479
39. Lewandowski D (2016) Perspektiven eines Open Web Index. Inf Wiss Prax 67:15–21. https://doi.org/10.1515/iwp-2016-0020
40. Lewandowski D (2014) Why we need an independent index of the web. In: König R, Rasch M (eds) Society of the query reader: reflections on web search. Institute of Network Culture, Amsterdam, pp 49–58
41. Blank D, Fuhr N, Henrich A, Mandl T, Rölleke T, Schütze H, Stein B (2011) Teaching IR: curricular considerations. In: Efthimiadis E (ed) Teaching and learning in information retrieval. Springer, Berlin, Heidelberg, New York, pp 31–46 https://doi.org/10.1007/978-3-642-22511-6_3
42. Markov I, de Rijke M (2019) What should we teach in information retrieval? ACM SIGIR Forum 52:19–39. https://doi.org/10.1145/3308774.3308780
43. Lewandowski D, Sünkler S, Schultheiß S (2020) Studies on search: designing meaningful IIR studies on commercial search engines. Datenbank Spektrum 20:5–15. https://doi.org/10.1007/s13222-020-00331-1