



Universiteit
Leiden
The Netherlands

The user perspective in professional information search

Verberne, S.; Vossen, P; Fokkens, A.

Citation

Verberne, S. (2022). The user perspective in professional information search. In P. Vossen & A. Fokkens (Eds.), *Creating a more transparent internet* (pp. 229-243). Cambridge University Press. Retrieved from <https://hdl.handle.net/1887/3736317>

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/3736317>

Note: To cite this publication please use the final published version (if applicable).

The User Perspective in Professional Information Search

Suzan Verberne

17.1 Introduction

Information retrieval (IR) is the research field that addresses the development, optimization, and evaluation of search engines and the study of how humans interact with search engines. Traditionally, search systems were investigated in the context of libraries and librarians (Harman et al., 2019), but the main focus of the field since the 1990s has been on Web search engines that serve an immense target audience (e.g., Google).

A large part of IR research is involved in the development of ranking optimization methods – a machine learning – driven task. But the field has never lost sight of the user perspective (Croft, 2019; Ingwersen and Järvelin, 2005). The human is central in the information search process because the user formulates queries, views the retrieved documents, judges their relevance, and decides when to stop searching (Maxwell, 2019). The user's perspective on information is important in the IR process in two different ways: (1) the user formulates a search query based on their own perspective of the task at hand and the required information; and (2) the user assesses the results returned by the search engine for relevance to their information need. Depending on the user's perspective, a document could be relevant or irrelevant to the entered query. Although the user acts from their personal perspective, the search engine only sees the user's interactions: the entered queries and the clicked documents. For example, consider a user entering the query "Rembrandt Leiden" in a Web search engine. The perspective of the user could be a historical one, searching information about Rembrandt van Rijn's youth in Leiden. It could also be a touristic perspective, searching for Rembrandt locations in Leiden that are worth visiting. Or maybe the user's goal is simply to navigate to the Web page of the restaurant Rembrandt in Leiden.

In this chapter, we use the term *perspective* in the context of IR to refer to all user aspects that lead to the formulated query and the assessment of the results: the user's interest, background information, current task context, and information need.

Because of the central role of the user in IR, it feels natural to take the individual user's perspective into account in the development of search engines. Current search engines use a form of personalization in the ranking of the search results, adapting the ranking for a small portion of the search results. This form of personalization could, for example, accommodate for local search, ranking search results from locations close to the user higher than results that are further away (Hannak et al., 2013). But the large majority of the search results in Web search engines are in fact user-independent: each user receives the same results without any personalization. The ranking is strongly effected by the popularity of pages – estimated by the clicks of other users (Joachims et al., 2005). This is effective for common Web search tasks but not for the highly specific search tasks performed in professional contexts.

Professional search is the searching carried out by experts for work purposes (Russell-Rose et al., 2018; Verberne et al., 2019a). Or, in the words of Russell-Rose et al., based on Tait (2014):

Professional search focuses on the work of paid professionals who are undertaking a work task that is predominately search-related and performed under a number of constraints such as budget and time. (Russell-Rose et al., 2018)

Professional search is a relatively small area of research in the IR field, but it has been recognized as an important application domain that is challenging because of the specific needs of professional search engine users (Russell-Rose et al., 2018; Salampasis et al., 2013; Verberne et al., 2018b, 2019b). Personalization of search results could potentially increase the effectiveness of professional search, but when developing these methods, we should be aware that professional users – more than users in generic Web search – need to be in control of the search process and must be able to trust the system to provide them with reliable information. Thus, transparency of the retrieval system is essential in this context.

This chapter addresses the user perspective in professional search. In Section 17.2 we introduce professional search as a research area. In Section 17.3 we discuss relevant work on personalization in information retrieval. We then summarize the recent studies addressing explainable search and recommendation in Section 17.4, and in Section 17.5 we give an outline for research directions in the near future, aiming at explainable professional search that makes the user perspective central in the search process.

17.2 Professional Search

Professional searchers, such as lawyers, information specialists, policy officers, architects, and scholars, need to process increasing amounts of documents to find relevant, complete, high-quality, work-related information (Bawden and Robinson, 2009; Sappelli, 2016).

In the common Web search paradigm, as implemented by search engines such as Google and DuckDuckGo, result ranking largely relies on popularity of Web pages: the more hyperlinks from popular pages link to a document, and the more often a document is clicked for a given query, the higher it is ranked in future searches (Joachims et al., 2005). For example, soon after Wikipedia became popular on the Web, Wikipedia pages started to end up on the top of the list on Google result pages for many queries – and this is still the case (DuckDuckGo, 2020). In day-to-day Web search, many users have the same information needs, and therefore popularity is a relevant ranking criterion.

A problem that arises when this popularity-driven paradigm is applied to work-related search is that the most popular documents are often not the most relevant documents for the individual user in their current search task. The differences between professional information search and generic Web search can be summarized in three important aspects:

- The search tasks of professionals are complex, that is, highly specific and typically recall-oriented: the searchers want to be sure that they have found all the relevant information (Kim et al., 2011; Mason, 2006);
- The searching is not limited to sending one query and clicking one result, but is often exploratory by nature (He et al., 2013), and includes browsing, analyzing (Makri et al., 2008), and re-finding previously used information (Sappelli et al., 2017);
- Each user has their own individual needs: not only interests, expertise, and information needs differ per user, but also the perceived relevance of retrieved documents (Sun et al., 2008). The search evolves on the searcher's own knowledge.

Because the information needs in professional search are highly specific and individual, the relevance of the results depends heavily on the *user perspective*. Therefore, the click data available from other users is limited and irrelevant (Huang et al., 2016). Hence, result ranking cannot depend on popularity. An alternative is to use the searcher's own history for improving the search ranking. To achieve that, a user profile must be created and utilized for personalized ranking (Micarelli and Sciarrone, 2004). This brings us to the topic of personalized IR.

17.3 Personalized IR

User profiling and personalization have been addressed extensively in IR research (Ghorab et al., 2013). Approaches to user profiling and personalization typically learn user preferences by collecting queries and clicked documents (Micarelli et al., 2007). A rich user profile can be learned by extracting prominent terms from the clicked documents and storing them in a term profile (Tang et al., 2010; Teevan et al., 2005). Often, the extracted terms are connected to an existing domain knowledge base, for example, a legal thesaurus or a medical ontology (Daoud et al., 2009; Speretta and Gauch, 2005). The term “profile” can then be used to better help the user find relevant information.

One way to do that is to re-order the results based on similarity to the user profile, where the documents that are in the interest field of the user are ranked higher (Micarelli and Sciarrone, 2004). For identifying which documents are relevant to the user profile, and which are not, it is sometimes necessary to perform *query disambiguation* (Tanudjaja and Mui, 2002). Queries are often short, and the user has a specific underlying intent in mind that is unknown to the search engine. A classic example is the query “java,” which can refer to either the island or the programming language. The user profile can help in deciding which of the two meanings is more of interest to the user.

The user profile can also be used for personalized query expansion (Zhou et al., 2012) – expanding the user with relevant terms based on the domain of the user, or query suggestion (Leung et al., 2008; Verberne et al., 2015) – showing terms to the user that are likely to be relevant additions to the query. For example, when I enter the query “search behaviour” in Google Scholar (see Figure 17.1), the first three results are about marine predator search behavior, visual search behavior in expert soccer goalkeepers, and job search behavior. Of course, as an IR researcher, I am interested in *information* search behavior, and adding the word “information” to my query improves the relevance of the search results. A query *suggestion* module could detect this and help me improve my query; a query *expansion* module could in the background compare my query with my user profile and add user-specific terms to my query automatically.

Although all research cited here reports an improvement of personalization over the non-personalized baseline, the actual implementation of personalization strategies in search environments is limited: on average, only 11.7% of Google Web Search results show differences due to personalization (Hannak et al., 2013; Hannák et al., 2017). This is because users are wary when it comes to personalization; they feel that their privacy is violated when the search engine uses their personal information. Privacy-preserving personalization is




Figure 17.1 Example of personal relevance in academic search (Google Scholar).
 (a) The query “search behaviour” gives results that are irrelevant for me; (b) Google’s query suggestion functionality suggests specifications of my query. The first suggestion is relevant to me; (c) If I search for “information search behaviour” I receive relevant results. Google and the Google logo are trademarks of Google LLC

therefore an important societal topic (Karwatzki et al., 2017; Mittelstadt, 2016). A crucial step in the development of privacy-secure systems is to make the system transparent and explainable to the user (Holzinger et al., 2017). This is further elaborated in the next section.

Transparency and explainability are even more important in work-related contexts than in the Web search context: professional users do not want to feel that they are losing control over the search process because the ranking of the search results is not stable or not predictable (Russell-Rose et al., 2018).

17.4 Explainable Recommendation and Search

In artificial intelligence, explainability is an important means to address issues with transparency and trust in black-box machine learning models (Adadi and Berrada, 2018). Search and recommender systems are ubiquitous in our daily lives, and it is considered important that users understand the results and recommendations they receive (Zhang et al., 2019). Recommender systems such as Amazon, Spotify, and Netflix recommend items to users that they are likely to appreciate (buy, listen, watch), without waiting for a user's query. Search systems provide information that is estimated to be relevant to a query entered by the user. In both types of systems, it can be valuable for the user to see why information was presented to them. In explainable *personalized* search, the user gets an answer to the question: "Why was this information presented to me," and even "Why should I trust the given information?"

Explainability in Recommender Systems

Early approaches to recommendation were inherently explainable because they were relatively straightforward, directly related to the content of the suggested items and their user ratings. Examples of explainable recommendations in these systems are "You have highly rated items that are similar to this item" and "Users who have similar ratings with you highly rated this item." With the rise of machine learning methods in recommender systems for a large audience (e.g., Netflix, Spotify), the implementation of explainability and its trade-off with accuracy became more challenging (Koren et al., 2009; Zhang et al., 2014). The trade-off is based on the discrepancy that machine learning systems, and in particular deep neural network approaches, are outperforming rule-based systems in the quality of recommended results, but they are much more difficult to explain to the user than simple rule-based systems.

In recent years, the use of knowledge graphs to facilitate explainability recommendation methods has been proposed (Ai et al., 2018; Wang et al., 2018). Knowledge graphs are flexible and can integrate heterogeneous information

types. Users and items are both modeled as nodes (entities) in the graph and the strengths of the relations between entities are used for recommending new entities to a user. Explanations can be generated in natural language to explain the relevance of a specific item to the user (Balog et al., 2019).

Explainability in Search Systems

Personalized search has in common with recommender systems that the user profile determines the relevance of a document. But as opposed to recommender systems, retrieval systems get an input query and need to retrieve documents that are relevant to that particular query.

In explainable search, the aim is that the user knows the capabilities and limitations of the search system, that they trust the system, and know how to intervene with the system if the results are not satisfactory (Zhang et al., 2019). A common form of explanation of the relevance of search results is the use of search snippets on the result page in which query terms have been marked in boldface. This markup (which is used in all Web search engines) indicates the topical relevance of a document for the user query. Search snippets are a basic example of explanations unified with the ranking model: Ranking models have term overlap metrics as central components, and term overlap is directly visualized in the snippets on the result page.

For relevance factors that, as opposed to term overlap, do not directly follow from the ranking model, a separate explanation engine is needed that generates explanations post hoc. This need has become more urgent as the state of the art in IR is now held by deep neural network models that do not use human-defined features for ranking the documents but abstracted document representations (Nogueira and Cho, 2019; Yang et al., 2019). Since 2019, efforts have been made to make features and their importance weights from neural retrieval models explicit and visualize these features as an explanation of why a document is relevant to the user query (Chios and Verberne, 2020; Fernando et al., 2019; Singh and Anand, 2019).

In the context of *professional* search, explainability is a novel research direction with no experimental results published yet at the time of writing. In the next section, we list suggestions for research in this direction.

17.5 Toward Explainable Professional Search

After having discussed professional search (Section 17.2), personalized search (Section 17.3), and explainable search (Section 17.4) in the previous three sections, we now bring these topics together in the next step for advancing

professional search: the development of explainable search methods in the professional context that allow for personalization without becoming a black box to the user.

According to Russell-Rose and MacFarlane (2020 p. 2), explainability in professional search has two criteria: (1) the ease with which the user's information need can be translated into a query (explainability of the query process), and (2) the degree to which the user's query returns the results expected and intended by the user (explainability of the search results). Explainability of the query process strongly relates to professional query interfaces, which often allow the user to build complex Boolean queries. Russell-Rose and MacFarlane recommend improving the explainability of query interfaces by providing real-time feedback on query effectiveness, allowing users to evaluate the contribution of individual query elements (Russell-Rose and MacFarlane, 2020; Russell-Rose and Shokraneh, 2020, p. 4).

In this chapter, we focus on the second criterion: the explainability of the search results.

Since professional search tasks are highly specific, result ranking cannot rely on the data of other users. Given this individual nature of professional search, personalization of professional search seems a logical step. However, when using user information in relevance ranking, it is important for users to have insight in to the data that is stored by the search engine (Xu et al., 2007) and to understand the influence of their personal data on the search results. Thus, professional search relies on explainable models in order to have the user trust the system and be in control of the search process.

The current state of research, discussed in the previous sections, gives way to two research directions for the near future:

1. Post hoc explanations added to the ranked lists of documents;
2. Graph-based personalized search, explicitly adding the individual user perspective to the searching and browsing process.

17.5.1 Explanations for Estimated Document Relevance

Just as snippets give an indication of topical relevance by highlighting query terms in text excerpts, other relevance factors could also be explicitly highlighted on the result page. These relevance factors could differ between domains. For example, users of a legal search engine (lawyers, legal scholars, legal professionals) consider document characteristics such as source authority, legal hierarchy, and whether the document is annotated to be important factors of relevance (Wiggers et al., 2018). Adding such metadata information of the retrieved documents to the result page is relatively straightforward; the next step would be to show indications of the weight that the ranking model

assigned to each relevance factor. This helps inform the user about which factors were taken into account for the ranking and how they were weighted. A paper by Chios and Verberne (2020) proposed a search engine result page on which the relative importance of query terms for the retrieved documents and the position of the most relevant passage in each document are shown. This was positively valued by the participants of a small-scale user study: they give significantly higher scores for the explainability and assessability (how well can the relevance of the retrieved documents be assessed) of the result page. This paper could be followed up by work addressing relevance factors in professional contexts.

17.5.2 Explainable Search Using a Personal Graph

A promising direction for explainable search is the use of graph models. Graphs are a natural and transparent means of representing knowledge (Chein and Mugnier, 2008). Knowledge graphs have been shown to be especially helpful in exploratory search tasks (Sarrafzadeh et al., 2014, 2016), which are common in professional work environments (He et al., 2013). Graphs can also be used for generating search explanations by explicitly describing the path between users and items in the graph. If we take academic search (Chiang et al., 2013; Salehi et al., 2015; Verberne et al., 2015, 2018b) as an example, we could generate explanations such as “this article is retrieved because you have previously read papers that cite it, and because you commonly read papers from this journal.”

Most previous works in graph-based search use an external knowledge graph covering all domain knowledge. Verberne (2018) and Balog and Kenter (2019) have both proposed storing personal knowledge graphs to enable personalized search. A personal knowledge graph is “a resource of structured information about entities personally related to its user, their attributes and the relations between them” (Balog and Kenter, 2019). Thus, the personal knowledge graph is a possible visualization of the user’s perspective in the search process. In the proposal by Verberne (2018), the personal knowledge graph is a professional graph representing the searching and browsing history of the user in the professional search engine. A graph representing the knowledge and interests of one user is much smaller than a graph representing the complete index of a search engine (Blanco and Lioma, 2012) and can be stored locally (client-side), if privacy regulations require it.

There are two main challenges associated with the idea of a personal graph for information search: automatically populating the personal graph from sparse user data and effectively utilizing the graph for effective information finding.

Future research with professional knowledge graphs should address the development of methods for these two aspects.

17.6 Conclusions

In this chapter, we have discussed the idea of the Perspective Web in the context of IR, and in particular information search for professional purposes. In their search for information, users act based on their own, personal perspective. User queries in Web search engines are often underspecified because much of the user context is implicit. The underspecification of user queries leads to ambiguity: Does the query term “search behavior” refer to *predator* search behavior, *job* search behavior, or *information* search behavior? The user knows, but the search engine does not. Modern Web search engines solve this by showing a diversity of perspectives to the user, hoping that a relevant result is among these. In the ranking of results in Web search engines, popularity is an important criterion: the more users have clicked on a Web page, the more often it shows up in the result list of other users.

Professional search tasks are user- and context-specific. This means that the user perspective plays an even larger role in the relevance of the returned results than in Web search; marine predator search behavior would be relevant from the perspective of a marine biologist, but not from the perspective of me as an IR researcher. Thus, ranking algorithms cannot use the popularity of search results as effectively as in Web search. This establishes the potential for personalization in professional search. At the same time, professional users want to be in control over the search process and need to be able to trust the search engine to provide them with correct and relevant information. This motivates the necessity to make the results retrieved by the professional search engine explainable to the user. We have discussed the state of the art in explainable recommendation and search and then proposed two possible research directions for explainable professional search: adding explanations to traditional result pages and developing a search paradigm that is centered around the user’s professional knowledge graph. This latter research direction could potentially lead to a true realization of the user perspective in information search.

There is one caveat to a search engine that centers around the user’s perspective, and that is the filter bubble effect (Nguyen et al., 2014): if the user profile is based on the user’s past behavior and the user profile is used to change the search results, the risk is that the user will dive deeper in directions that confirm their own beliefs (perspective), ignoring the results that contradict them. This is one of the reasons why user control is important: the user needs to see at any time what the influence of their user profile is on the results they see.

The explainable interface needs to include visual information on the user perspective itself. In my vision, the graph visualization proposed in Section 17.5.2 would become a kaleidoscope where a different perspective changes the view of the data.

References

- Adadi, A., and Berrada, M. 2018. Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, **6**, 52138–52160.
- Ai, Q., Azizi, V., Chen, X., and Zhang, Y. 2018. Learning heterogeneous knowledge base embeddings for explainable recommendation. *Algorithms*, **11**(9), 137.
- Balog, K., and Kenter, T. 2019. Personal knowledge graphs: A research agenda. Pages 217–220 of: *Proceedings of the 2019 ACM SIGIR International Conference on Theory of Information Retrieval*. ICTIR '19. New York: ACM.
- Balog, K., Radlinski, F., and Arakelyan, S. 2019. Transparent, scrutable and explainable user models for personalized recommendation. Pages 265–274 of: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '19. New York: ACM.
- Bawden, D., and Robinson, L. 2009. The dark side of information: Overload, anxiety and other paradoxes and pathologies. *Journal of Information Science*, **35**(2), 180–191.
- Blanco, R., and Lioma, C. 2012. Graph-based term weighting for information retrieval. *Information Retrieval*, **15**(1), 54–92.
- Chein, M., and Mugnier, M.-L. 2008. *Graph-based knowledge representation: computational foundations of conceptual graphs*. London: Springer Science & Business Media.
- Chiang, M.-F., Liou, J.-J., Wang, J.-L., Peng, W.-C., and Shan, M.-K. 2013. Exploring heterogeneous information networks and random walk with restart for academic search. *Knowledge and Information Systems*, **36**(1), 59–82.
- Chios, I., and Verberne, S. 2020. Helping results assessment by adding explainable elements to the deep relevance matching model. Pages 1–9 of *Proceedings of the 3rd International Workshop on ExplainAble Recommendation and Search*. Xi'an: EARS '20.
- Croft, W. B. 2019. The importance of interaction for information retrieval. Pages 1–2 of: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '19. New York: ACM.
- Daoud, M., Lechani, L.-T., and Boughanem, M. 2009. Towards a graph-based user profile modeling for a session-based personalized search. *Knowledge and Information Systems*, **21**(3), 365–398.
- DuckDuckGo 2020. *Information retrieval at DuckDuckGo*. <https://duckduckgo.com/?t=ffab&q=information+retrieval&ia=web>.
- Fernando, Z. T., Singh, J., and Anand, A. 2019. A study on the interpretability of neural retrieval models using DeepSHAP. Pages 1005–1008 of: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. Paris: ACM.

- Ghorab, M. R., Zhou, D., O'Connor, A., and Wade, V. 2013. Personalised information retrieval: Survey and classification. *User Modeling and User-Adapted Interaction*, **23**(4), 381–443.
- Hannák, A., Sapiezynski, P., Khaki, A. M., Lazer, D., Mislove, A., and Wilson, C. 2017. Measuring personalization of Web search. *arXiv preprint arXiv:1706.05011*.
- Hannak, A., Sapiezynski, P., Molavi Kakhki, A., Krishnamurthy, B., Lazer, D., Mislove, A., and Wilson, C. 2013. Measuring personalization of Web search. Pages 527–538 of: *Proceedings of the 22nd International Conference on World Wide Web. WWW '13*. New York: ACM.
- Harman, D. 2019. Information retrieval: The early years. *Foundations and Trends® in Information Retrieval*, **13**(5), 425–577.
- He, J., Bron, M., and de Vries, A. P. 2013. Characterizing stages of a multi-session complex search task through direct and indirect query modifications. Pages 897–900 of: *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*. Dublin: ACM.
- Holzinger, A., Biemann, C., Pattichis, C. S., and Kell, D. B. 2017. What do we need to build explainable AI systems for the medical domain? *arXiv preprint arXiv:1712.09923*.
- Huang, Z., Cautis, B., Cheng, R., and Zheng, Y. 2016. KB-enabled query recommendation for long-tail queries. Pages 2107–2112 of: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. Indianapolis: ACM.
- Ingwersen, P., and Järvelin, K. 2005. *The turn: Integration of information seeking and retrieval in context (The information retrieval series), Vol. 18*. Dordrecht: Springer.
- Joachims, T., Granka, L., Pan, B., Hembrooke, H., and Gay, G. 2005. Accurately interpreting clickthrough data as implicit feedback. Pages 154–161 of: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York: ACM.
- Karwatzki, S., Dytynko, O., Trenz, M., and Veit, D. 2017. Beyond the personalization–privacy paradox: Privacy valuation, transparency features, and service personalization. *Journal of Management Information Systems*, **34**(2), 369–400.
- Kim, Y., Seo, J., and Croft, W. B. 2011. Automatic Boolean query suggestion for professional search. Pages 825–834 of: *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '11. New York: ACM.
- Koren, Y., Bell, R., and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer*, **42**(8), 30–37.
- Leung, K. W.-T., Ng, W., and Lee, D. L. 2008. Personalized concept-based clustering of search engine queries. *IEEE Transactions on Knowledge and Data Engineering*, **20**(11), 1505–1518.
- Makri, S., Blandford, A., and Cox, A. L. 2008. Investigating the information-seeking behaviour of academic lawyers: From Ellis's model to design. *Information Processing & Management*, **44**(2), 613–634.
- Mason, D. 2006. Legal information retrieval study–Lexis professional and Westlaw UK. *Legal Information Management*, **6**(4), 246–250.

- Maxwell, D. M. 2019. Modelling search and stopping in interactive information retrieval. Ph.D. thesis, University of Glasgow.
- Micarelli, A., Gasparetti, F., Sciarrone, F., and Gauch, S. 2007. Personalized search on the World Wide Web. *The Adaptive Web*, **4321**, 195–230.
- Micarelli, A., and Sciarrone, F. 2004. Anatomy and empirical evaluation of an adaptive Web-based information filtering system. *User Modeling and User-Adapted Interaction*, **14**(2), 159–200.
- Mittelstadt, B. 2016. Auditing for transparency in content personalization systems. *International Journal of Communication*, **10**(June), 4991–5002.
- Nguyen, T. T., Hui, P.-M., Harper, F. M., Terveen, L., and Konstan, J. A. 2014. Exploring the filter bubble: The effect of using recommender systems on content diversity. Pages 677–686 of: *Proceedings of the 23rd International Conference on World Wide Web*. Seoul: ACM.
- Nogueira, R., and Cho, K. 2019. Passage re-ranking with BERT. *arXiv preprint arXiv:1901.04085*.
- Russell-Rose, T., Chamberlain, J., and Azzopardi, L. 2018. Information retrieval in the workplace: A comparison of professional search practices. *Information Processing & Management*, **54**(6), 1042–1057.
- Russell-Rose, T., and MacFarlane, A. 2020. Towards explainability in professional search. *Proceedings of the 3rd International Workshop on Explainable Recommendation and Search*. EARS '20.
- Russell-Rose, T., and Shokraneh, F. 2020. Designing the structured search experience: Rethinking the query-builder paradigm. *Weave: Journal of Library User Experience*, **3**(1).
- Salampasis, M., Fuhr, N., Hanbury, A., Lupu, M., Larsen, B., and Strindberg, H. 2013. Integrating IR technologies for professional search. Pages 882–885 of: Serdyukov, P., Braslavski, P., Kuznetsov, S. O., Kamps, J., Rüger, S., Agichtein, E., Segalovich, I., and Yilmaz, E. (eds), *Advances in information retrieval*. Berlin: Springer.
- Salehi, S., Du, J. T., and Ashman, H. 2015. Examining personalization in academic Web search. Pages 103–111 of: *Proceedings of the 26th ACM Conference on Hypertext & Social Media*. Guzelyurt: ACM.
- Sappelli, M. 2016. Knowledge work in context. User centered knowledge worker support. Ph.D. thesis, Radboud University Nijmegen.
- Sappelli, M., Verberne, S., and Kraaij, W. 2017. Evaluation of context-aware recommendation systems for information re-finding. *Journal of the Association for Information Science and Technology*, **68**(4), 895–910.
- Sarrafzadeh, B., Vechtomova, O., and Jokic, V. 2014. Exploring knowledge graphs for exploratory search. Pages 135–144 of: *Proceedings of the 5th Information Interaction in Context Symposium on - IliX '14*. Regensburg: ACM.
- Sarrafzadeh, B., Vtyurina, A., Lank, E., and Vechtomova, O. 2016. Knowledge graphs versus hierarchies. Pages 91–100 of: *Proceedings of the 2016 ACM on Conference on Human Information Interaction and Retrieval – CHIIR '16*. Carrboro: ACM.
- Singh, J., and Anand, A. 2019. EXS: Explainable search using local model agnostic interpretability. Pages 770–773 of: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. Melbourne: ACM.

- Speretta, M., and Gauch, S. 2005. Personalized search based on user search histories. Pages 622–628 of: *Web Intelligence, 2005. Proceedings: The 2005 IEEE/WIC/ACM International Conference on Web Intelligence*.
- Sun, Y., Li, H., Council, I. G., Huang, J., Lee, W.-C., and Giles, C. L. 2008. Personalized ranking for digital libraries based on log analysis. Pages 133–140 of: *Proceedings of the 10th ACM workshop on Web Information and Data Management*. Napa Valley: ACM.
- Tait, J. I. 2014. An introduction to professional search. Goos, G., Hartmanis, J., and van Leeuwen, J. (eds.) Pages 1–5 of: *Professional search in the modern world*. Springer.
- Tang, J., Yao, L., Zhang, D., and Zhang, J. 2010. A combination approach to Web user profiling. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 5(1), 2.
- Tanudjaja, F., and Mui, L. 2002. Persona: A contextualized and personalized Web search. Pages 1232–1240 of: *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*. Big Island: IEEE.
- Teevan, J., Dumais, S. T., and Horvitz, E. 2005. Personalizing search via automated analysis of interests and activities. Pages 449–456 of: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. Salvador: ACM.
- Verberne, S. 2018. Explainable IR for personalizing professional search. Pages 35–42 of: *Joint Proceedings of the First International Workshop on Professional Search (ProfS2018); the Second Workshop on Knowledge Graphs and Semantics for Text Retrieval, Analysis, and Understanding (KG4IR); and the International Workshop on Data Search*. Ann Arbor: SEARCH'18.
- Verberne, S., He, J., Kruschwitz, U., Larsen, B., Russell-Rose, T., and de Vries, A. P. 2018b. First international workshop on professional search (ProfS2018). Pages 1431–1434 of: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. SIGIR '18. New York: ACM.
- Verberne, S., He, J., Kruschwitz, U., Wiggers, G., Larsen, B., Russell-Rose, T., and de Vries, A. P. 2019a. First international workshop on professional search. Pages 153–162 of: *ACM SIGIR Forum*, Vol. 52. New York: ACM.
- Verberne, S., He, J., Wiggers, G., Russell-Rose, T., Kruschwitz, U., and de Vries, A. P. 2019b. Information search in a professional context-exploring a collection of professional search tasks. *arXiv preprint arXiv:1905.04577*.
- Verberne, S., Kraaij, W., and de Vries, A. P. 2018a. Author-topic profiles for academic search. *arXiv preprint arXiv:1804.11131*.
- Verberne, S., Sappelli, M., Kalervo, J., Järvelin, K., and Kraaij, W. 2015. User simulations for interactive search: Evaluating personalized query suggestion. Pages 678–690 of: *Advances in information retrieval*, Vol. 9022. Vienna: Springer International.
- Wang, H., Zhang, F., Wang, J., Zhao, M., Li, W., Xie, X., and Guo, M. 2018. Ripplet: Propagating user preferences on the knowledge graph for recommender systems. Pages 417–426 of: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. Torino: ACM.
- Wiggers, G., Verberne, S., and Zwenne, G.-J. 2018. Exploration of intrinsic relevance judgments by legal professionals in information retrieval systems. Pages 5–8 of: *Proceedings of the 17th Dutch-Belgian Information Retrieval Workshop*. Leiden: arXiv.

- Xu, Y., Wang, K., Zhang, B., and Chen, Z. 2007. Privacy-enhancing personalized Web search. Pages 591–600 of: *Proceedings of the 16th International Conference on World Wide Web*. Banff: ACM.
- Yang, W., Zhang, H., and Lin, J. 2019. Simple applications of bert for ad hoc document retrieval. *arXiv preprint arXiv:1903.10972*.
- Zhang, Y., Lai, G., Zhang, M., Zhang, Y., Liu, Y., and Ma, S. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. Pages 83–92 of: *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval*. Gold Coast: ACM.
- Zhang, Y., Mao, J., and Ai, Q. 2019. SIGIR 2019 tutorial on explainable recommendation and search. Pages 1417–1418 of: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. Paris: ACM.
- Zhou, D., Lawless, S., and Wade, V. 2012. Improving search via personalized query expansion using social media. *Information Retrieval*, **15**(3–4), 218–242.