Implementing Knowledge Gain Measurement in Real Search Environments

Rafael Tavares da Silva[1], **Sean Wolfgand Matsui Siqueira**[1] **(Orientador)**
Marcelo Tibau[1] **(Co-Orientador)**

[1] Department of Applied Informatics – Federal University of the State of Rio de Janeiro Av. Pasteur, 458, Urca – 22.290-240 – Rio de Janeiro – RJ – Brazil

rafaeltavares.silva@edu.unirio.br,

{sean,marcelo.tibau}@uniriotec.br

Abstract. *The operationalization of learning metrics in real search environments remains an open challenge in the Searching as Learning (SaL) paradigm. While behavioral proxies offer scalability, they capture conceptual change only indirectly; structured assessments provide more direct evidence but often compromise ecological validity. The Degree of Knowledge Gain (DKG) metric addresses this tension by combining Shannon entropy with semantic similarity between queries and clicked documents to model the progressive reduction of uncertainty during search. This paper reports on two technological artifacts developed to embed DKG computation into real-world search workflows, within the scope of the CNPq project 3C-BPA: Comportamento de busca, Complexidade da informação e pensamento Crítico na Busca como um Processo de Aprendizagem. A standalone search engine prototype established the feasibility of real-time DKG computation but exposed limitations in ecological validity and operational sustainability. These were addressed by a Chrome browser extension that estimates the metric unobtrusively while users interact with their preferred search engines. To assess the extension's applicability, an experiment was conducted combining pre- and post-tests with the Concurrent Think-Aloud (CTA) protocol and automated interaction logging. Preliminary results indicate that DKG is sensitive to variation in search strategy use as participants who engaged in systematic query reformulation and multi-source evaluation achieved stronger knowledge gains, while those exhibiting disorientation and limited cognitive regulation showed more modest outcomes. Beyond its empirical contributions, the study illustrates how undergraduate research participation can play a substantive role in advancing the development and application of formal learning metrics in information science.*

1. Introduction

Web search has increasingly been understood not merely as an act of information retrieval, but as a knowledge construction process, consistent with the Searching as Learning (SaL) perspective [Marchionini 2006, Vakkari 2016, Tibau et al. 2022]. A persistent challenge in this domain concerns how to assess what users actually learn during a search session. Behavioral indicators such as clicks and query reformulations are scalable and easy to capture, yet they offer only indirect evidence of conceptual change [Xu et al. 2020, Gritz et al. 2021, Tibau et al. 2019]. Structured assessments, in turn, can provide more direct evidence of learning, though often at the cost of ecological validity [Schmuckler 2001]. The Degree of Knowledge Gain (DKG) metric was proposed to address this tension by modeling learning as a reduction of uncertainty grounded in Shannon entropy, while also incorporating semantic similarity between queries and clicked documents [Tibau 2024]. Although prior work has demonstrated the metric's interpretability and alignment with learning indicators, an important next step is to examine how it can be operationalized in real search environments [Tibau 2024, Tibau et al. 2022, Tibau et al. 2023].

This paper reports on that effort within the scope of the CNPq project *3C-BPA: Comportamento de busca, Complexidade da informação e pensamento Crítico na Busca como um Processo de Aprendizagem*. Specifically, it focuses on the development of technological artifacts designed to embed DKG computation into actual search workflows. The work was carried out primarily through an undergraduate research project, in which the student implemented and tested two complementary artifacts, a prototype search engine with built-in DKG computation and a browser plug-in capable of estimating the metric in real time as users interact with their preferred search engines. Together, these developments advanced the practical viability of the metric and created an opportunity for formative research training at the undergraduate level.

To investigate the plug-in's applicability, the authors designed an experiment combining pre- and post-tests with the Concurrent Think-Aloud (CTA) protocol [Kelley et al. 2015]. Participants completed tasks of varying complexity, spanning introductory science topics to more applied themes such as prosthetics. This setup enabled examination of both the behavior of DKG as an automated measure and the search strategies participants employed throughout the sessions.

Beyond its technical contribution, this work is positioned within the field of Information Systems by addressing the interplay between people, processes, and technology in search-based learning environments. From the perspective of people, the study examines how users employ different search strategies, regulate their behavior, and construct knowledge while interacting with search systems. From the perspective of processes, it investigates how knowledge acquisition unfolds during search sessions, including query reformulation, source evaluation, and decision-making during learning-oriented tasks. From the perspective of technology, it advances artifacts capable of capturing these dynamics automatically in real search environments.

From a broader research perspective, this work also aligns with the Brazilian Grand Challenges in Information Systems (GranDSI-BR 2016–2026), particularly with respect to the sociotechnical, complexity, and open-world dimensions of IS. It contributes to a sociotechnical vision by integrating behavioral, cognitive, and technological perspectives; addresses the complexity of Information Systems through a formal and operationalizable metric that combines probabilistic modeling, interaction data, and semantic relationships; and situates itself in the open-world context by investigating learning processes in real web search environments characterized by heterogeneous information sources and evolving interaction patterns.

The remainder of this paper is organized as follows. Section 2 reviews related work on Searching as Learning and existing approaches to measuring knowledge gain. Section 3 presents the principles of the DKG metric and the artifacts developed for its operationalization. Section 4 describes the experimental design, including the role of the undergraduate researcher in the study's setup, execution, and analysis. Section 5 reports the preliminary findings and their implications. Finally, Section 6 reflects on the metric's potential, its limitations, and directions for future work, highlighting both methodological contributions and the formative role of undergraduate participation.

2. Related Work

Research in SaL has consistently held that search is not merely a means of locating information, but an active process of knowledge construction [Urgo and Arguello 2022]. A central challenge within this perspective is determining the extent to which learning actually takes place during a search session. Approaches to this problem fall broadly into two categories: those that rely on direct measures of user knowledge, such as pre/post-tests, open-ended responses, and self-reports, and those that infer learning indirectly from interaction data, such as behavioral logs and query traces [Tibau 2024].

A substantial body of work has explored indirect approaches, often treating behavioral signals as assumed proxies for learning. Early studies [Chi et al. 2016] and subsequent research [Xu et al. 2020, Otto et al. 2021, Yu et al. 2021, Gritz et al. 2021, Otto et al. 2022, El Zein and da Costa Pereira 2022] examined signals such as query reformulations, document sequences, and interaction patterns as indicators of evolving knowledge states. More recent contributions have extended this line of inquiry through knowledge-graph models designed to approximate conceptual growth during exploratory tasks [El Zein and Da Costa Pereira 2023, Liu et al. 2023], as well as multimodal methods, including eye-tracking, to analyze reading behavior and its relationship with knowledge acquisition [Gritz et al. 2024]. Other studies have drawn attention to the role of task-related and contextual factors, such as query refinement behavior and motivational attributes, in shaping how users learn through search [Câmara 2024, Liu et al. 2025].

Overall, these findings reinforce the view that search systems can do more than retrieve information: they can also support, stimulate, and measure learning. Even so, existing behavioral proxies carry important limitations. They tend to treat learning as an emergent consequence of interaction rather than modeling its underlying cognitive structure, and while they offer scalability, they often capture knowledge acquisition in coarse terms that may not adequately reflect qualitative differences across tasks of varying complexity.

The DKG metric was proposed to address these limitations by combining Shannon entropy with semantic similarity between queries and clicked documents to formalize the progressive refinement of knowledge during search [Tibau 2024]. Unlike conventional behavioral proxies, DKG models learning as a reduction of uncertainty grounded in observable interaction patterns, making it a promising candidate for operationalization in environments designed to retrieve information, and to evaluate or stimulate learning. The present work advances this agenda by exploring how DKG can be embedded into real-world search settings through practical technological artifacts and experimental use within an educational research context.

3. The DKG Metric and the Development of Artifacts

The DKG metric provides a formal model for quantifying learning during search. It is grounded in Shannon's Entropy (H), which represents uncertainty as a probability distribution over possible outcomes [Prieto-Guerrero and Espinosa-Paredes 2019]. In the context of search, this uncertainty corresponds to the user's knowledge state before encountering new information; as learning takes place, entropy is expected to decrease. To capture this process more meaningfully, DKG also incorporates semantic similarity between queries and clicked documents, so that knowledge gain reflects both novelty and conceptual refinement.

The qualitative basis of DKG draws on the formulation of additive information gain [Milne 2012], according to which each new information element β contributes incrementally to what is already known from α:

$$i(\alpha \land \beta, \Gamma) = i(\alpha, \Gamma) + i(\beta, \Gamma \cup \alpha) \tag{1}$$

where Γ denotes the evolving knowledge base. This iterative view of knowledge acquisition informed the design of artifacts capable of estimating DKG in real time.

The formal expression of DKG is given in [Tibau 2024] as:

$$\deg_{KG} = \left(1 - \sum_{i=1}^{n} p_i \log \frac{m_i}{p_i}\right) \times 0.01 \tag{2}$$

where:

- p_i is the probability of detecting a clicked document at position i, estimated from empirical click-through data Dupret-Piwowarski-2008;
- m_i is the observation factor, which combines the click indicator (c_i) with the Jaccard similarity between successive queries.

Thus, m_i captures the joint influence of user behavior, such as clicks and query reformulations, and system behavior, such as ranking position, while p_i models the likelihood that these interactions occur. The scaling factor preserves the interpretability of the metric.

Compared with simpler proxies such as click counts or dwell time, DKG combines rank-weighted detectability (p_i) and inter-query similarity (through m_i), yielding a single interpretable estimate of knowledge refinement throughout a search session.

3.1. First Artifact: Search Engine Prototype

The first implementation effort consisted of a software prototype that simulated a conventional search engine while embedding DKG computation into its internal logic. Its purpose was to automate the calculation of the metric during live search sessions, thereby enabling the immediate tracking of users' knowledge gain. Developed over approximately two months, the prototype was employed in preliminary experiments and played an important role in demonstrating the feasibility of the approach. However, important limitations became evident since the retrieved results did not always align with participants' search intentions, which negatively affected both user experience and the validity of the collected data. In addition, the external API used to provide search results generated high token costs per session, making broader experimentation impractical. Although the prototype fulfilled its role as a proof of concept, these constraints made clear the need for a more sustainable and ecologically valid solution.

To address the limitations of the prototype, a second artifact was developed in the form of a Google Chrome plug-in. Unlike the standalone search engine, the extension operated within users' natural search environment, allowing DKG to be computed unobtrusively in the background while users continued to rely on their preferred search engines. This design improved ecological validity and reduced the intrusiveness of the data collection process.

From a technical perspective, the extension stored session data in Supabase, which provided structured persistence through an adequate free-tier infrastructure. This eliminated the prohibitive costs associated with API-based retrieval and enabled the consistent recording of query, click, ranking, and semantic similarity data. Methodologically, the extension also made it possible to conduct experiments under more naturalistic conditions, improving reliability and user experience.

Implementation at a glance:

- Stack: Chrome MV3 extension, background logger, and Supabase storage.
- Signals logged: queries, clicks, ranks, inter-query Jaccard similarity, and timestamps (used in the computation of m_i and p_i).
- Privacy: local pseudonyms were adopted, and no page content was stored, only interaction metadata.

Experiments with five participants provided initial evidence of the plug-in's viability. The decision to work with five users is consistent with a long-standing tradition in usability research, according to which small samples are often sufficient to identify core interaction issues and reach observational saturation [Lu 2025]. In comparison with the first artifact, the extension produced more relevant results, required no maintenance costs, and supported a more scalable approach to experimentation.

3.2. Comparison and Contribution

The transition from the standalone search engine to the Chrome extension represents an important step in the operationalization of the DKG metric. Although the first artifact was limited in practical terms, it established the feasibility of computing DKG in real time. The second artifact advanced this effort by offering a more sustainable and methodologically robust solution, better suited to real search environments. Developed within the context of an undergraduate research project, these artifacts show how a theoretical model such as DKG can be translated into tools for use in naturalistic search settings. By allowing more scalable and ecologically valid experimentation, they constitute both a technical contribution and a methodological bridge between the formalization of knowledge gain and practical applications of Searching as Learning.[1]

4. Methodology

The study was designed to evaluate the DKG metric in real search scenarios through a combination of automatic logging and qualitative inquiry. Five university volunteers, recruited through convenience sampling and with self-reported prior familiarity

[1] The artifacts are available at `https://doi.org/10.5281/zenodo.17203699`

with web search, completed three tasks of increasing complexity under controlled conditions. Ethical approval was obtained from the institutional review board (CAAE 82881624.0.0000.5285), and all participants provided informed consent. The sessions were anonymized, recorded, and subsequently transcribed. An important aspect of the study was the involvement of the undergraduate student supported by the project *Comportamento de busca, Complexidade da informação e pensamento Crítico na Busca como um Processo de Aprendizagem (3C-BPA)*. His contribution was central across multiple stages of the experiment. First, he prepared the experimental environment by configuring the logging plug-in, setting up the online platform for the pre- and post-tests, and ensuring the integration of these components. During the sessions, he also assisted participants by clarifying task instructions and helping ensure that the Concurrent Think-Aloud (CTA) protocol [Kelley et al. 2015] was followed consistently.

Beyond data collection, the student also contributed to the analytical process by preparing transcripts that support the preliminary qualitative findings and by advancing the ongoing content analysis. This stage involves coding participant interactions using the Online Information Searching Strategies (OISS) framework [Reisoğlu et al. 2019] and the ESKiP Taxonomy of Query State [Tibau et al. 2019], both of which enable a systematic mapping of search strategies and query transitions.

On the quantitative side, the study combined pre- and post-test scores with automatically computed DKG values derived from query and click logs. These different sources of evidence were then integrated through mixed-methods triangulation, linking external learning outcomes with DKG measurements and coded strategies[2]. This combination illustrates both the methodological strength of the study and the formative role of undergraduate research in advancing the development and application of the DKG metric.

5. Findings

The preliminary analysis of the transcripts revealed a range of search behaviors shaped both by participants' prior familiarity with search engines and by their ability to regulate the search process in real time. Procedural strategies associated with Trial-and-Error were especially frequent. Participants often reformulated queries, tested alternative keywords, and explored different websites when their initial attempts were unsuccessful. User 4 exemplified this pattern by persistently refining queries after receiving unhelpful results, demonstrating both determination and iterative problem-solving.

Metacognitive strategies were more evident among participants who achieved higher post-test gains. Purposeful Thinking, for instance, appeared in behaviors such as narrowing the scope of a query or conducting targeted searches within specific websites. User 3, for example, consistently evaluated and cross-checked multiple sources, weighing their credibility and synthesizing information across sites. At the same time, indicators of disorientation were also observed. User 2, for instance, showed frustration when initial searches failed, at times abandoning attempts quickly or reacting negatively to unsuccessful outcomes. These patterns suggest that limited cognitive regulation may hinder effective learning through search.

[2]The anonymized transcripts are available at `https://doi.org/10.5281/zenodo.17195699` and the dataset is available at `https://doi.org/10.5281/zenodo.17202855`.

From the perspective of the ESKiP taxonomy, participants relied most heavily on Specialization (SC) and Word Substitution (WS). These reformulation moves enabled them to narrow broad queries or adjust vocabulary to better match the language expected in relevant documents. User 5 illustrates this pattern, beginning with more general queries and progressively adding detail to obtain more precise results. By contrast, Generalization (GE) moves were rare, suggesting that participants more often reduced than expanded the scope of their queries when facing obstacles. Across sessions, higher DKG values tended to co-occur with more frequent purposeful reformulations and cross-source evaluation, whereas moments marked by disorientation were associated with flatter DKG trajectories. Although still preliminary, this pattern was observed consistently across all sessions.

Considered jointly, these findings suggest that the DKG metric is sensitive to variations in strategy use. While the automated measures indicate that learning takes place, the transcript analysis helps explain why outcomes differ across participants. Persistent reformulation and structured source evaluation, as seen in Users 4 and 3, were associated with stronger gains, whereas frustration and less systematic behavior, as observed in User 2, corresponded to more modest improvements.

Beyond their descriptive value, these findings also have implications for the design of Information Systems intended to support learning through search. The observed association between purposeful query reformulation, source evaluation, and higher knowledge gain suggests that search systems and learning platforms could incorporate mechanisms to detect and encourage such behaviors in real time. For instance, adaptive interfaces might offer prompts for query refinement, highlight diverse or higher-quality sources, or signal patterns associated with disorientation. In this sense, the DKG metric may function as an evaluative measure as well as a foundation for feedback-driven systems that actively support knowledge acquisition during search.

6. Conclusion

This study contributes to the ongoing effort to operationalize the Degree of Knowledge Gain (DKG) metric in real search environments through the development and testing of two artifacts: a prototype search engine and a browser extension. Together, these artifacts demonstrate how a theoretical model of knowledge gain can be embedded into practical tools for use in authentic search settings. The transition from the standalone prototype to the Chrome extension represented an important methodological advance, as it increased ecological validity, reduced operational costs, and enabled experimentation under more naturalistic conditions. The mixed-methods experiment provided preliminary evidence that DKG is sensitive to differences in search strategy use. Participants who engaged more systematically in query reformulation and source evaluation tended to achieve higher levels of knowledge gain, as reflected in both pre/post-test comparisons and DKG scores, whereas disorientation and limited cognitive regulation were associated with more modest gains. Taken together, these results reinforce the potential of DKG as both a diagnostic and an evaluative tool in Searching as Learning.

Another important aspect of this work was the formative role of undergraduate research. The student supported by the 3C-BPA project contributed substantially to the design, implementation, and execution of the study, including the development and refinement of the artifacts, the preparation of the experimental environment, support during data collection, and the ongoing content analysis. His involvement strengthened the practical viability of the DKG metric and illustrates how undergraduate research can simultaneously stimulate innovation and provide meaningful training in empirical methods. Although the findings remain preliminary, they point to the promise of DKG as a scalable and interpretable measure of knowledge acquisition during search.

Beyond their methodological and empirical contributions, the results also point to broader implications for Information Systems. The proposed artifacts may support a range of institutions that rely on search-intensive knowledge work, from educational settings and digital learning platforms to organizational environments in which professionals must acquire, verify, and apply information during task execution. In such contexts, estimating knowledge gain from interaction traces may help improve learning workflows, inform the design of adaptive support mechanisms, and provide decision makers with evidence about how users engage with search systems.

These possibilities open concrete avenues for the integration of DKG across Information Systems contexts. In educational environments, the metric may support intelligent tutoring systems and digital learning platforms by providing continuous estimates of knowledge acquisition during search-based activities. In search engines and information access systems more broadly, DKG could inform the personalization of results or interface elements based on inferred learning progress.

Nevertheless, the findings reported in this study should be interpreted as exploratory, reflecting an initial proof-of-concept for the operationalization of the DKG metric in real search environments. The experimental scope remains limited, particularly due to the small sample size (n = 5) and the controlled nature of the study, which constrain the generalizability of the results. While this design is consistent with established practices in early-stage usability and exploratory research, further validation is required to assess the robustness and scalability of the metric across broader contexts.

In future work, the authors intend to report a full, artifact-centered study with an expanded participant pool, greater task variety, and more comprehensive reliability analyses, including detailed examinations of DKG components. Future investigations will also explore longitudinal designs to better capture knowledge acquisition over time and its relation to sustained search behaviors. In addition, the authors intend to triangulate CTA with complementary protocols (e.g., concurrent/retrospective hybrids and design-thinking probes [Kelley et al. 2015] to strengthen process validity.

References

Câmara, A. (2024). *Designing Search-as-Learning Systems*. Ph.d. thesis, Delft University of Technology. ⟨DOI: 10.4233/uuid:0fe3a6bb-1bc1-40e2-86b0-ec3d3aef9c77⟩, ⟨ISBN: 978-94-6384-569-4⟩.

Chi, Y., Han, S., He, D., and Meng, R. (2016). Exploring knowledge learning in collaborative information seeking process. In *SAL@SIGIR*.

El Zein, D. and da Costa Pereira, C. (2022). User's knowledge and information needs in information retrieval evaluation. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*, UMAP '22, page 170–178, New York, NY, USA. Association for Computing Machinery.

El Zein, D. and Da Costa Pereira, C. (2023). The evolution of user knowledge during search-as-learning sessions: A benchmark and baseline. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval*, CHIIR '23, page 454–458, New York, NY, USA. Association for Computing Machinery.

Gritz, W., Hoppe, A., and Ewerth, R. (2021). On the impact of features and classifiers for measuring knowledge gain during web search-a case study.

Gritz, W., Hoppe, A., and Ewerth, R. (2024). On the influence of reading sequences on knowledge gain during web search. In Goharian, N., Tonellotto, N., He, Y., Lipani, A., McDonald, G., Macdonald, C., and Ounis, I., editors, *Advances in Information Retrieval*, pages 364–373, Cham. Springer Nature Switzerland.

Kelley, T., Capobianco, B., and Kaluf, K. (2015). Concurrent think-aloud protocols to assess elementary design students. *International Journal of Technology & Design Education*, 25(4):521 – 540.

Liu, C., Song, X., and Hansen, P. (2023). Characterising users' task completion process in learning-related tasks:a search pace model. *Journal of Information Science*, 49(6):1462–1480.

Liu, Y., Qin, C., Ma, X., Chen, J., He, H., and Mao, J. (2025). Characterising exploratory search tasks: Evidence from different fields. *Journal of Information Science*, 0(0):01655515251330611.

Lu, J. 2025). A systematic literature review of usability definitions and psychometric properties of instruments in the field of learning design and technology. *Journal of Research on Technology in Education*, 0(0):1–22.

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

Milne, P. (2012). Probability as a Measure of Information Added. *Journal of Logic, Language and Information*, 21(2):163–188.

Otto, C., Rokicki, M., Pardi, G., Gritz, W., Hienert, D., Yu, R., von Hoyer, J., Hoppe, A., Dietze, S., Holtz, P., Kammerer, Y., and Ewerth, R. (2022). SaL-Lightning Dataset: Search and Eye Gaze Behavior, Resource Interactions and Knowledge Gain during Web Search. In *ACM SIGIR Conference on Human Information Interaction and Retrieval*, pages 347–352, Regensburg Germany. ACM.

Otto, C., Yu, R., Pardi, G., von Hoyer, J., Rokicki, M., Hoppe, A., Holtz, P., Kammerer, Y., Dietze, S., and Ewerth, R. (2021). Predicting knowledge gain during web search based on multimedia resource consumption. In Roll, I., McNamara, D., Sosnovsky, S., Luckin, R., and Dimitrova, V., editors, *Artificial Intelligence in Education*, pages 318–330, Cham. Springer International Publishing.

Prieto-Guerrero, A. and Espinosa-Paredes, G. (2019). 7 - nonlinear signal processing methods: Dr estimation and nonlinear stability indicators. In Prieto-Guerrero, A. and Espinosa-Paredes, G., editors, *Linear and Non-Linear Stability Analysis in Boiling Water Reactors*, Woodhead Publishing Series in Energy, pages 315 – 398. Woodhead Publishing.

Reisoğlu, I., Çebi, A., and Bahçekapılı, T. (2019). Online information searching behaviours: examining the impact of task complexity, information searching experience, and cognitive style. *Interactive Learning Environments*, 0(0):1–18.

Schmuckler, M. A. (2001). What is ecological validity? a dimensional analysis. *Infancy*, 2(4):419–436.

Tibau, M. (2024). *Quantifying Knowledge Gain in Online Searches: The DKG Metric*. Tese (doutorado em informática), Universidade Federal do Estado do Rio de Janeiro (UNIRIO), Rio de Janeiro. Programa de Pós-Graduação em Informática.

Tibau, M., Siqueira, S. W. M., and Nunes, B. P. (2022). The impact of non-verbalization in think-aloud: Understanding knowledge gain indicators considering think-aloud web searches. In *Proceedings of the 33rd ACM Conference on Hypertext and Social Media*, HT '22, page 107–120, New York, NY, USA. Association for Computing Machinery.

Tibau, M., Siqueira, S. W. M., and Nunes, B. P. (2023). Accounting for the knowledge gained during a web search: An empirical study on learning transfer indicators. *Library & Information Science Research*, 45(1):101222.

Tibau, M., Siqueira, S. W. M., Nunes, B. P., Nurmikko-Fuller, T., and Manrique, R. F. (2019). Using query reformulation to compare learning behaviors in web search engines. In *2019 IEEE 19th International Conference on Advanced Learning Technologies (ICALT)*, volume 2161-377X, pages 219–223.

Urgo, K. and Arguello, J. (2022). Learning assessments in search-as-learning: A survey of prior work and opportunities for future research. *Information Processing & Management*, 59(2):102821.

Vakkari, P. (2016). Searching as learning: A systematization based on literature. In *Journal of Information Science, 42(1)*, pages 7–18.

Xu, L., Zhou, X., and Gadiraju, U. (2020). How does team composition affect knowledge gain of users in collaborative web search? In *Proceedings of the 31st ACM Conference on Hypertext and Social Media*, HT '20, page 91–100, New York, NY, USA. Association for Computing Machinery.

Yu, R., Tang, R., Rokicki, M., Gadiraju, U., and Dietze, S. (2021). Topic-independent modeling of user knowledge in informational search sessions. *Information Retrieval Journal*, 24(3):240–268.