Tutorial on Task-Based Search and Assistance

Chirag Shah
University of Washington
Seattle, Washington
chirags@uw.edu

Ryen W. White
Microsoft Research
Redmond, Washington
ryenw@microsoft.com

ABSTRACT

While great strides are made in the field of search and recommendation, there are still challenges and opportunities to address information access issues that involve solving tasks and accomplishing goals for a wide variety of users. Specifically, we lack intelligent systems that can detect not only the request an individual is making (what), but also understand and utilize the intention (why) and strategies (how) while providing information. Many scholars in the fields of information retrieval, recommender systems, productivity (especially in task management and time management), and artificial intelligence have recognized the importance of extracting and understanding people's tasks and the intentions behind performing those tasks in order to serve them better. However, we are still struggling to support them in task completion, e.g., in search and assistance, it has been challenging to move beyond single-query or single-turn interactions. The proliferation of intelligent agents has opened up new modalities for interacting with information, but these agents will need to be able to work more intelligently in understanding the context and helping the users at task level. This tutorial will introduce the attendees to the issues of detecting, understanding, and using task and task-related information in an information episode (with or without active searching). Specifically, it will cover several recent theories, models, and methods that show how to represent tasks and use behavioral data to extract task information. It will then show how this knowledge or model could contribute to addressing emerging retrieval and recommendation problems.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Relevance assessment**; *Personalization.*

KEYWORDS

Tasks, Search and recommendation systems, Intelligent assistants

ACM Reference Format:
Chirag Shah and Ryen W. White. 2020. Tutorial on Task-Based Search and Assistance. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20), July 25–30, 2020, Virtual Event, China.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3397271.3401422

SIGIR '20, July 25–30, 2020, Virtual Event, China

ACM ISBN 978-1-4503-8016-4/20/07.
https://doi.org/10.1145/3397271.3401422

1 MOTIVATION

"A wealth of information creates a poverty of attention." – this statement by renowned economist and psychologist Herbert A. Simon from 1971 is perhaps even more fitting today as we are drenched in the sea of information on one hand, but lack tools and support for going beyond retrieving relevant information and meet higher level goals of solving problems and making decisions. While great strides are made in the field of search and recommendation, there are still challenges and opportunities to address information access issues that involve solving tasks and accomplishing goals for a wide variety of users. Specifically, we lack intelligent systems that can detect not only the request an individual is making (what), but also understand and utilize the intention (why) and strategies (how) while providing information. Many scholars in information retrieval (IR) and beyond have recognized the importance of extracting and understanding user's task and the intention behind doing that task in order to serve them better [26, 27, 35], but we are still struggling to get out of single-query or single-turn interactions. The proliferation of intelligent agents has opened up new modality for interacting with information, but these agents will need to be able to work more intelligently in understanding the context and helping the people at the task level. Several events have taken place over the past few years around this topic (e.g., [4, 16, 19]) (PAIR Workshop, 2020; WICRS Workshop, 2020), indicating continual and increased interest by scholars, but also highlighting the need to do more in this important area.

As outlined in Section 3, while there has been growing interest in this area as reflected in the scholarly outcomes and events over the last few years, an educational effort such as a tutorial has been missing. We aim to address this by delivering a half-day tutorial on task-based search and intelligent assistance that informs and inspires the attendees to go beyond addressing single-turn retrieval and recommendation problems to connecting these processes to searcher context and the task at hand.

2 OBJECTIVES

Following are the learning objectives to be addressed in this tutorial:

- Recognize **situations** where task knowledge can be useful in fulfilling a person's need and in helping them complete their current task.
- Discuss various **methods** to extract task information (both from log data and in real-time), including the topic (what), people's strategies and processes (how), and their intentions (why).
- Describe various types of **task support** offered by search engines, digital assistants, and task management applications.
- Outline different evaluation **metrics** for assessing the performance of intelligent systems.
- Identify **challenges and opportunities** in making progress in this area and the role that the IR community can play.

3 RELEVANCE TO THE IR COMMUNITY

This section covers how the proposed topic is relevant to the IR community, along with how it has been positioned and studied over the years in the literature. The section also provides pointers to some of the relevant events in the recent past in this space.

3.1 Related Research

Information search and retrieval systems are a significant component of advanced intelligent assistance. People's interactions with search systems are often motivated by tasks that emerge from evolving, continuous problematic situations [5, 6]. Understanding a searcher's intention and task for asking a question could provide a general framework for conceptualizing different contexts and situations of information needs that drive people to seek information using different information systems. Although search and retrieval systems have evolved, they became more advanced in suggesting useful queries, and providing personalized search results, especially in fact-finding and navigational search tasks [37], their algorithmic processes are still limited in retrieving information supporting complex and intellectually challenging tasks [16]. Thus, often, current information systems fail to provide information relevant to achieve searchers' complex task goals, or are unable to suggest useful queries to exclude or execute next in order to retrieve Web pages containing useful information for searchers' tasks.

Rooted in cognitive perspective, task-based approach in information seeking and retrieval emerged within the interactive IR (IIR) community with studies conducted by Vakkari [35] and Ingwersen and Järvelin [18], which consider tasks in the design of IR systems to find out for what purposes the system is used [28] and thus provides implications for IR system design to personalize information according to the task at hand. Based on a series of empirical works, Vakkari [35] developed a general framework of task-based information searching which consists of three stages: *pre-focus*, *focus formulation*, and *post-focus*.

Many studies have investigated and identified various aspects of task and especially interactive nature of search tasks. Bates' [3] berrypicking model showed the interactive process of searching. The aim of task-based studies is to investigate the relationships between task characteristics and information seeking behaviors by recognizing and understanding the nature of different tasks and goals and designing IR systems which can support accomplishment of a variety of such tasks and goals. Here, the task is a multi-level information seeking process in which people need information to achieve a goal to fulfill the task (e.g., [10, 28, 29, 34]). Many existing task models (e.g., [11, 20, 22]) have investigated and identified searchers' tasks as static and overarching goals that motivate search actions.

Different characteristics or facets of tasks [22] influence people's interaction with intelligent systems, for example, when a searcher searches for information using a search engine [24]. Search tasks are influenced by the work task or everyday life task that drives them to seek information or are associated with a problematic situation [9]. Also, accomplishing more complex tasks require more complex actions that are manifested throughout the session because complex tasks take longer to complete or require more queries.

Apart from task, existing studies in IR segment information seeking behaviors into various levels of explicit and implicit signals. While performing tasks, searchers' actions are also driven by intentions and can be well-defined or ill-defined [18]. These studies have indicated that there is a close association between searchers' performance of a task and the information need, the search strategies employed, and the assessment of document relevance and utility.

Thus, an understanding of searchers' information seeking goals at the task level is essential to improve intelligent systems because different searchers have different needs and intentions; they face different problems in different situations. The usefulness or relevance of the information for a searcher may vary based on that searcher's specific situation or the context in which the information is needed or used. As the information searching session progresses, a searcher interacts with new information, which may change their state of knowledge, thus changing their needs and usefulness of documents. As a result, without considering the evolving nature of dynamic tasks, the recommendations provided by a search system may not be useful to a searcher in their current situation.

In recent years, advanced machine learning and deep learning model-based recommendations have made enormous progress and somewhat succeeded in solving these problems [12]. By extracting meaningful latent factors from highly diverse and complex heterogeneous data, deep learning models such as matrix factorization, auto-encoders, memory networks, neural networks with collaborative filtering can achieve better recommendations [17, 23]. Furthermore, incorporating structure-based or feature-based knowledge graphs with a sequential learning or alternative learning approaches into a recommendation system have been shown to improve its performance [36]. Reinforcement learning-based approaches leverage feedback generated from people's continuously evolving interactions with the system, along with historical data, to generate recommendations [41].

Beyond search, tasks pervade almost every aspect of our daily work and personal lives [2]. They involve different activities, have different constraints, and take different amounts of time to complete. Some tasks can be completed quickly, while others take much longer, sometimes spanning several days or weeks. Task management applications such as Microsoft To Do, Google Tasks, and Todoist help people track their pending and completed tasks. Studies have found that users of these and similar systems would benefit from assistance with many aspects of task management, especially task planning [8]. Scheduling and prioritizing tasks are both challenging [27]. There has been some recent progress in task intelligence, in areas such as discovering digital assistant capabilities [38], estimating how long tasks will take to complete [39], and automatically tracking task status over time [40].

3.2 Related Events

Several workshops have been held in the space of task-based IR. Many of these workshops have focused on search interactions, searcher intents and tasks in information search, including CHI 2012 workshop on *End-user Interactions with Intelligent Systems* [33] organized by Simone Stumpf, Margaret Burnett, Volkmar Pipek, and Weng-Keen Wong, and the *Second Strategic Workshop on Information Retrieval in Lorne (SWIRL)* [1]. In 2012, Birger Larsen, Christina Lioma, and Arjen de Vries organized *Task-based and Aggregated Search Workshop* [21], which focused on the challenges of task-based and aggregated search such as, the mismatch between

search interface and specialized task-based functionalities, the lack of homogeneous systems to support different tasks, and so on. One of the significant contributions of this early workshop was that it identified how and to what extent a domain-specific search and recommendation systems could be developed to support task level activities. Participants also discussed how a retrieval system should be modified in order to provide better support for task-based search. In the same year, Nicholas Belkin, Charles Clarke, Ning Gao, Jaap Kamps, and Jussi Karlgren organized the *SIGIR 2012 Workshop on "Entertain Me" Supporting Complex Search Tasks* [7]. The interactive workshop hosted by ACM and SIGIR brought together researchers from different backgrounds in fostering potential solutions to problems faced by searchers with complex information needs. Aiming to support searchers during their entire search sessions when interactively solving a complex task, the workshop explored many aspects of interactive information systems such as complex search episodes, queries, exploratory search, understanding of search context, and finally, how to incorporate task and searcher context into an information system.

In another SIGIR workshop on *Modeling User Behavior for Information Retrieval Evaluation* [13], participants examined ways to model search intent based on queries. They also identified problems with the use of queries as a proxy for search intent and brainstormed better solutions. In the first and second workshops on *Supporting Complex Search Tasks* organized by Maria Gade, Mark Hall, Hugo Huurdeman, Jaap Kamps, Marijn Koolen, Mette Skov, Elaine Toms, David Walsh [14], and Nick Belkin, Toine Bogers, Jaap Kamps, Diane Kelly, Marjin Koolen, and Emine Yilmaz [4] respectively in 2015 and 2017 also prepared to initiate an interdisciplinary dialogues among researchers from information retrieval, information behavior, human-computer interaction, and computer science addressing many task-related open research questions. Participants tackled issues related to six aspects of information seeking – context, tasks, heterogeneous sources and search process, UI and UX, and evaluation of systems. The workshops were helpful in fostering new collaborations among different communities to address these issues.

A more recent workshop hosted at WSDM 2018 and organized by Rishabh Mehrotra, Ahmed H. Awadallah, and Emine Yilmaz on *Learning from User Interactions* [25] focused on task-based intelligent systems, more specifically at six related topics – user needs and tasks understanding, user modeling and personalization, metrics and evaluation, user interaction processes and context, intelligent interface design and applications. The workshop attracted participants from IR, human factors, ubiquitous computing, data mining, and other related domains. The *Task Intelligence Workshop* [16], organized by Ahmed H. Awadallah, Mark Sanderson, Cathal Gurring, and Ryen White focused on various topic related to tasks in the context of system development including areas such search assistance, personalization, and recommendation.

4 FORMAT AND SCHEDULE

This will be a half-day tutorial with the following tentative schedule:

- **Hour-1:** Introduction and discussion with examples and situations where task knowledge is important in understanding and addressing user needs. These will come from various domains including intelligent assistants, e-commerce, and health. This will be followed by discussing some of the recent and seminal works on frameworks and models around task identification and representation, within and beyond the IR community.
- **Hour-2:** Deeper dive into current methods for using behavioral and other signals to extract aspects of a task. These aspects include the nature of the search (e.g., exploratory, fact-finding) and the nature of the outcome expected (e.g., intellectual, product-based), among other things.
- **Hour-3:** Discussion on various methods for evaluating task-based search and assistance systems. Specifically, thinking through questions such as the following:

(1) How do we measure user satisfaction?
(2) How do we determine task progress and task completion?
(3) How do we assess the performance of systems supporting zero-query retrieval and/or recommendations?

Importantly, to help ground the discussion, we will frame the tutorial content in terms of real-world task support currently offered by search systems, intelligent assistants, and task management applications. For instance, we will discuss how task knowledge could be essential and integrative in being able to develop a multi-turn conversational search system.

5 MATERIAL

The following material will be made available to tutorial attendees to help them make the most of their participation and experience.

- Presentation slides
- Annotated bibliography on task-based search and assistance

These items will be shared through GitHub, with the link provided to attendees during the tutorial and shared with the community afterwards.

6 ABOUT THE PRESENTERS

Chirag Shah is an Associate Professor in the Information School (iSchool) at the University of Washington (UW) in Seattle. Before UW, he was a faculty at Rutgers University. His research interests include studies of interactive information retrieval/seeking, trying to understand the task a person is doing and providing proactive recommendations. Dr. Shah received his Ph.D. in Information Science from University of North Carolina (UNC) at Chapel Hill. He directs the InfoSeeking Lab where he investigates issues related to information seeking, human-computer interaction (HCI), and fairness in machine learning. Shah has authored/edited two books on Collaborative Information Seeking (CIS) [15, 30], and a book on Social Information Seeking (SIS) [31]. He was a guest editor for the IEEE Computer Special Issue on CIS published in March 2014. His new textbook on Data Science was recently released [32]. He has taught undergraduate and graduate courses in IR, HCI, and Data Science at University of North Carolina (UNC) Chapel Hill, Rutgers University, and University of Washington. He has also taught several courses and tutorial on topics related to IR at different international places, including at SIGIR and WSDM conferences, Russian Summer School on Information Retrieval (RuSSIR), and Asian Summer School in Information Access (ASSIA).

Ryen White is a Partner Research Manager at Microsoft Research AI, where he leads a world-class team of scientists and engineers. In recent roles, Ryen led the applied science organization for Microsoft Cortana and he was chief scientist for Microsoft

Health. Ryen's research has historically been focused on understanding search interaction and on developing tools to help people search more effectively. He received his Ph.D. in Computer Science from University of Glasgow, United Kingdom. Ryen has published hundreds of conference papers and journal articles in web search and many other areas. He was identified as the "Center of the SIGIR Universe" (most central author in the co-authorship graph) in the 40 years of the ACM SIGIR conference. Ryen has received many best-paper awards in conferences, journals and elsewhere, including at SIGIR (2007, 2010, 2013), ACM CIKM (2014, 2015), ACM SIGCHI (2011), and in JASIST (2010). His doctoral research received the British Computer Society's (BCS) Distinguished Dissertation Award for the best Computer Science Ph.D. dissertation in the United Kingdom in 2004/2005. In 2014, Ryen received the Microsoft BCS/BCS IRSG Karen Spärck Jones Award for contributions to information retrieval. He is a Fellow of the BCS.

7 ACKNOWLEDGEMENTS

This work in part is supported through a National Science Foundation (NSF) grant IIS-2017134 and an Amazon Research Award.

REFERENCES

[1] James Allan, Bruce Croft, Alistair Moffat, and Mark Sanderson. 2012. Frontiers, challenges, and opportunities for information retrieval: Report from SWIRL 2012 the second strategic workshop on information retrieval in Lorne. In *ACM SIGIR Forum*, Vol. 46. 2–32.

[2] David Allen. 2015. *Getting Things Done: The Art of Stress-free Productivity.* Penguin.

[3] Marcia J Bates et al. 1989. The design of browsing and berrypicking techniques for the online search interface. *Online Review* 13, 5 (1989), 407–424.

[4] Nicholas Belkin, Toine Bogers, Jaap Kamps, Diane Kelly, Marijn Koolen, and Emine Yilmaz. 2017. Second workshop on supporting complex search tasks. In *Proceedings of the ACM SIGIR CHIIR Conference on Conference Human Information Interaction and Retrieval*. 433–435.

[5] Nicholas J Belkin. 2008. Some (what) grand challenges for information retrieval. In *ACM SIGIR Forum*, Vol. 42. 47–54.

[6] Nicholas J Belkin. 2016. People, Interacting with Information. In *ACM SIGIR Forum*, Vol. 49. 13–27.

[7] Nicholas J Belkin, Charles LA Clarke, Ning Gao, Jaap Kamps, and Jussi Karlgren. 2012. Report on the SIGIR workshop on" entertain me" supporting complex search tasks. In *ACM SIGIR Forum*, Vol. 45. 51–59.

[8] Victoria Bellotti, Brinda Dalal, Nathaniel Good, Peter Flynn, Daniel G Bobrow, and Nicolas Ducheneaut. 2004. What a to-do: Studies of task management towards the design of a personal task list manager. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*. 735–742.

[9] Katriina Byström and Preben Hansen. 2005. Conceptual framework for tasks in information studies. *Journal of the American Society for Information Science and Technology* 56, 10 (2005), 1050–1061.

[10] Katriina Byström and Kalervo Järvelin. 1995. Task complexity affects information seeking and use. *Information Processing & Management* 31, 2 (1995), 191–213.

[11] Robert Capra, Jaime Arguello, Heather O'Brien, Yuan Li, and Bogeum Choi. 2018. The effects of manipulating task determinability on search behaviors and outcomes. In *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*. 445–454.

[12] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishi Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the First Workshop on Deep Learning for Recommender Systems*. 7–10.

[13] Charles LA Clarke, Luanne Freund, Mark D Smucker, and Emine Yilmaz. 2013. Report on the SIGIR 2013 workshop on modeling user behavior for information retrieval evaluation (MUBE 2013). In *ACM SIGIR Forum*, Vol. 47. 84–95.

[14] Maria Gäde, Mark M Hall, Hugo Huurdeman, Jaap Kamps, Marijn Koolen, Mette Skove, Elaine Toms, and David Walsh. 2015. Report on the first workshop on supporting complex search tasks. In *ACM SIGIR Forum*, Vol. 49. 50–56.

[15] Preben Hansen, Chirag Shah, and Clause-Peter Klas (Eds.). 2015. *Collaborative Information Seeking - Best Practices, New Domains, New Thoughts.* Springer. http://www.springer.com/us/book/9783319185415

[16] Ahmed Hassan Awadallah, Cathal Gurrin, Mark Sanderson, and Ryen W White. 2019. Task Intelligence Workshop@ WSDM 2019. In *Proceedings of the ACM WSDM Conference on Web Search and Data Mining*. 848–849.

[17] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the International Conference on the World Wide Web*. 173–182.

[18] Peter Ingwersen and Kalervo Järvelin. 2005. *The Turn: Integration of Information Seeking and Retrieval in Context.* Springer.

[19] Diane Kelly, Jaime Arguello, and Robert Capra. 2013. NSF workshop on task-based information search systems. In *ACM SIGIR Forum*, Vol. 47. 116–127.

[20] Diane Kelly, Jaime Arguello, Ashlee Edwards, and Wan-ching Wu. 2015. Development and evaluation of search tasks for IIR experiments using a cognitive complexity framework. In *Proceedings of the ACM SIGIR ICTIR International Conference on the Theory of Information Retrieval*. ACM, 101–110.

[21] Birger Larsen, Christina Lioma, and Arjen de Vries. 2012. Report on TBAS 2012: workshop on task-based and aggregated search. In *ACM SIGIR Forum*, Vol. 46. 71–77.

[22] Yuelin Li and Nicholas J Belkin. 2008. A faceted approach to conceptualizing tasks in information seeking. *Information Processing & Management* 44, 6 (2008), 1822–1837.

[23] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1754–1763.

[24] Jiqun Liu, Matthew Mitsui, Nicholas J Belkin, and Chirag Shah. 2019. Task, information seeking Intentions, and user behavior: Toward a multi-level understanding of web search. In *Proceedings of the ACM SIGIR CHIIR Conference on Human Information Interaction and Retrieval*. 123–132.

[25] Rishabh Mehrotra, Ahmed Hassan Awadallah, and Emine Yilmaz. 2018. Learnir: Wsdm 2018 workshop on learning from user interactions. In *Proceedings of the ACM WSDM Conference on Web Search and Data Mining*. 797–798.

[26] Matthew Mitsui and Chirag Shah. 2019. Bridging gaps: Predicting user and task characteristics from partial user information. In *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*. 415–424.

[27] Karen Myers, Pauline Berry, Jim Blythe, Ken Conley, Melinda Gervasio, Deborah L McGuinness, David Morley, Avi Pfeffer, Martha Pollack, and Milind Tambe. 2007. An intelligent personal assistant for task and time management. *AI Magazine* 28, 2 (2007), 47–47.

[28] Miamaria Saastamoinen and Kalervo Järvelin. 2017. Search task features in work tasks of varying types and complexity. *Journal of the Association for Information Science and Technology* 68, 5 (2017), 1111–1123.

[29] Reijo Savolainen. 2012. Expectancy-value beliefs and information needs as motivators for task-based information seeking. *Journal of Documentation* 68, 4 (2012), 492–511.

[30] Chirag Shah. 2012. *Collaborative Information Seeking: The Art and Science of Making the Whole Greater than the Sum of All (The Information Retrieval Series).* Springer. http://www.springer.com/computer/database+management+{&}+information+retrieval/book/978-3-642-28812-8

[31] Chirag Shah. 2017. *Social Information Seeking.* Springer.

[32] Chirag Shah. 2020. *A Hands-on Introduction to Data Science.* Cambridge University Press.

[33] Simone Stumpf, Margaret Burnett, Volkmar Pipek, and Weng-Keen Wong. 2012. End-user interactions with intelligent and autonomous systems. In *Proceedings of the ACM SIGCHI Extended Abstracts on Human Factors in Computing Systems*. 2755–2758.

[34] Pertti Vakkari. 1999. Task complexity, problem structure and information actions: Integrating studies on information seeking and retrieval. *Information Processing & Management* 35, 6 (1999), 819–837.

[35] Pertti Vakkari. 2001. A theory of the task-based information retrieval process: a summary and generalisation of a longitudinal study. *Journal of Documentation* 57, 1 (2001).

[36] Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. 2018. Ripplenet: Propagating user preferences on the knowledge graph for recommender systems. In *Proceedings of the ACM CIKM Conference on Information and Knowledge Management*. ACM, 417–426.

[37] Ryen W White. 2016. *Interactions with Search Systems.* Cambridge University Press.

[38] Ryen W White. 2018. Skill discovery in virtual assistants. *Commun. ACM* 61, 11 (2018), 106–113.

[39] Ryen W White and Ahmed Hassan Awadallah. 2019. Task duration estimation. In *Proceedings of the ACM WSDM Conference on Web Search and Data Mining*. 636–644.

[40] Ryen W White, Ahmed Hassan Awadallah, and Robert Sim. 2019. Task completion detection: A study in the context of intelligent systems. In *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*. 405–414.

[41] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. 2018. DRN: A deep reinforcement learning framework for news recommendation. In *Proceedings of the International Conference on the World Wide Web*. 167–176.