

Conversational Information Seeking: Theory and Application

Jeffrey Dalton
University of Glasgow
United Kingdom
jeff.dalton@glasgow.ac.uk

Sophie Fischer
University of Glasgow
United Kingdom
sophie.fischer@glasgow.ac.uk

Paul Owoicho
University of Glasgow
United Kingdom
p.owoicho.1@glasgow.ac.uk

Filip Radlinski
Google
United Kingdom
filiprad@google.com

Federico Rosetto
University of Glasgow
United Kingdom
federico.rosetto@glasgow.ac.uk

Johanne R. Trippas
RMIT University
Australia
johanne.trippas@rmit.edu.au

Hamed Zamani
University of Massachusetts Amherst
United States
zamani@cs.umass.edu

ABSTRACT

Conversational information seeking (CIS) involves interaction sequences between one or more users and an information system. Interactions in CIS are primarily based on natural language dialogue, while they may include other types of interactions, such as click, touch, and body gestures. CIS recently attracted significant attention and advancements continue to be made. This tutorial follows the content of the recent Conversational Information Seeking book [33] authored by several of the tutorial presenters. The tutorial aims to be an introduction to CIS for newcomers to CIS in addition to the recent advanced topics and state-of-the-art approaches for students and researchers with moderate knowledge of the topic. A significant part of the tutorial is dedicated to hands-on experiences based on toolkits developed by the presenters for conversational passage retrieval and multi-modal task-oriented dialogues. The outcomes of this tutorial include theoretical and practical knowledge, including a forum to meet researchers interested in CIS.

CCS CONCEPTS

• **Information systems** → **Information retrieval**; • **Computing methodologies** → **Discourse, dialogue and pragmatics**; • **Human-centered computing** → **Natural language interfaces**.

KEYWORDS

Conversational search; conversational question answering; conversational recommendation; dialogue systems; chatbots; interactive information retrieval

ACM Reference Format:

Jeffrey Dalton, Sophie Fischer, Paul Owoicho, Filip Radlinski, Federico Rosetto, Johanne R. Trippas, and Hamed Zamani. 2022. Conversational Information Seeking: Theory and Application. In *Proceedings of the 45th*

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR '22, July 11–15, 2022, Madrid, Spain

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8732-3/22/07.

<https://doi.org/10.1145/3477495.3532678>

International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22), July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3477495.3532678>

1 MOTIVATION

While CIS has roots in early information retrieval (IR) research, recent advances in automatic voice recognition and conversational agents created increasing interest in this area in IR as well as natural language processing and related fields. CIS was recognized as an emerging research area in the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018) [9].¹ Since then, it has grown into a mainstream research sub-field in major IR conferences and continues to evolve rapidly. The areas in CIS now include but are not limited to conversational search conceptualization (e.g., [5, 12, 19]), effective conversational query re-writing (e.g., [28]), generating and selecting clarifying questions (e.g., [30, 32]), conversational preference elicitation (e.g., [18, 35]), and understanding user interactions with spoken conversational systems (e.g., [25, 26]). The growing body of work in this area has been supplemented by an increasing number of recent seminars (e.g., [3]), workshops (e.g., [4, 6, 8, 23]), tutorials (e.g., [1, 13, 15, 16]), shared-tasks (e.g., [10, 20]), and datasets (e.g., [2, 7, 18, 21, 24, 27, 31]).

The purpose of this tutorial is to (1) introduce the topic for newcomers in the field, (2) present a big picture of the recent advances and their connections, (3) highlight the current open problems and directions for future explorations, and (4) suggest practical hands-on applications for CIS. The hands-on part provides experience using publicly available tools and resources. This includes Macaw from CIIR/Microsoft and tools from the GRILL Lab such as the interactive TREC CAsT toolkit (iCAsT), the OpenGRILL research platform used in the Amazon Alexa Taskbot challenge, and the Task Multimodal Agent Dialogue (Task MAD) [22] Wizard-of-Oz toolkit for data collection. Describing, but more importantly, using these systems provides theoretical and practical experience with new and emerging models and systems.

¹<https://sites.google.com/view/swirl3/>

2 OBJECTIVES

The tutorial aims to provide a broad overview of CIS, its basic concepts, and how these concepts interact. Our goal is to provide hands-on starting points for participants to experiment with CIS. The combined knowledge of the presenters and their enthusiasm for CIS will be a perfect starting point for newcomers, including updated techniques for intermediaries. Lastly, this tutorial acts as a forum for researchers interested in CIS.

3 RELEVANCE TO SIGIR

CIS is a core research program within IR. Papers related to conversational search, question answering, and recommender systems have been published in recent SIGIR conferences. In addition, the Conversational Assistance Track (CAST) has been organized by TREC since 2019. All of these demonstrate that there is significant interest in this topic from the SIGIR community.

4 RELATED TUTORIALS

A number of tutorials on the topic have been presented in recent years. They include:

- Recent Advances in Conversational Information Retrieval at SIGIR 2020 by Gao et al. [16]
- Interactive Information Retrieval: Models, Algorithms, and Evaluation at SIGIR 2021 by Zhai [34]
- Conversational Recommendation Systems at WSDM 2021 by Fu et al. [14]
- Conversational Recommendation: Formulation, Methods, and Evaluation at RecSys 2021 by Lei et al. [17]
- Conversational Information Seeking: Theory and Evaluation at CHIIR 2022 by Aliannejadi and Trippas [1]

Unlike the previous tutorials, the tutorial consists of a hands-on component that results in higher engagement and interactivity. In addition, the topics presented in the tutorial are well aligned with the CIS manuscript [33] which gives the audience extra reading material if they are interested in any part of the tutorial. Moreover, the progress in CIS has been rapidly evolving in the last few years, and the tutorial covers new developments in the field. Therefore, we believe that the tutorial has value even for those who attended the previous tutorials.

5 FORMAT AND DETAILED OUTLINE

Even though multiple tutorials appeared in previous IR-related venues, the tutorial is the first to include significant hands-on content with open-source tools on the topic. The tutorial targets the newcomers to the field in addition to researchers and students who have intermediate familiarity with CIS. Table 1 provides a tentative schedule showing the eight major parts.

Part 1: Introduction, Applications, and Historical Context. We start with introductory material to introduce the sub-field of CIS to the students and researchers who are new to the field. Therefore, we spend the first 45 minutes of the tutorial defining CIS, introducing its applications, and connecting with early seminal work related to CIS (e.g., I³R or THOMAS). This part of the tutorial is aligned with the first two chapters of the CIS manuscript [33] in addition to its Appendix A.

Part 2-3: Conversational Interfaces and Result Presentation, and Evaluating CIS Systems.

We continue by looking at CIS from the user perspective. We introduce different interfaces used in conversational systems (e.g., speech-only and multi-modal interfaces) and their unique properties. We also discuss result presentation using each of these interfaces and how users consume information. We further introduce approaches for evaluating CIS systems and explore how they are connected to user satisfaction. This part of the tutorial (a total of 75 minutes) is aligned with Chapters 3 and 7 in the CIS manuscript [33].

Part 4-6: Understanding Conversational Interactions, Conversational Response Ranking and Generation, and Mixed-Initiative Interactions.

We move towards methods for developing CIS systems and the core models in language understanding, ranking algorithms, ranking, and elements of initiative. We start with conversation understanding by introducing concepts related to topic shifting, coreference and ellipsis resolution. We then discuss various methods for ranking responses in conversational search, conversational question answering, and conversational recommendation systems. We also briefly discuss response generation, focusing on conversational question answering tasks. We then introduce types of mixed-initiative interactions, such as clarification and preference elicitation. We spend approximately 90 minutes on this part which is aligned with Chapters 4-6 from the CIS manuscript [33].

Part 7: Hands-on with CIS Platforms. The next part of the tutorial is hands-on. We provide platforms for real-time experimentation during the tutorial, so the audience can become familiar with the development of CIS systems. In more detail, our hands-on part would be based on the following tools:

OpenGRILL this provides hands-on experience using a version of the GRILLBot system developed for the Alexa Taskbot Challenge.² OpenGRILL supports hybrid CIS and task conversations incorporating conversational preference elicitation, conversational task search, and conversational question answering. It demonstrates how a CIS research prototype can be integrated with multiple live assistant frameworks, including Google’s DialogFlow to support real-world interactive experiments. Participants would gain hands-on experience with the latest cloud tools including, Docker, Kubernetes, gRPC, and a distributed architecture designed for diverse and scalable CIS use cases and real-world experimentation.

Interactive CAST (iCAST) Tools provides an interactive and hands-on experience with newly developed interactive CAST tools. This uses Colab and Jupyter Notebook to interact with cloud endpoints to support research on conversational search developed for TREC CAST 2021 [11]. This lets participants build and use a prototype CAST system in 30 minutes.

Task Multimodal Agent Dialogue (TaskMAD) [22] is an open-source platform for performing rich multimodal Wizard-of-Oz CIS experiments. It has been used in the Alexa Prize Challenge and at multiple universities for research and teaching. It supports both conversational search as well as hybrid task-oriented use cases. This would be a short introduction of approximately 15 minutes.

²<https://www.amazon.science/alex-prize/taskbot-challenge>

Table 1: Tentative schedule of the tutorial.

#	Time	Topic	Presenter	Associated Material and Chapters
1	45 min	Introduction, Applications, and Historical Context	Zamani	CIS manuscript, Ch. 1–2 & Appendix A
2	30 min	Conversational Interfaces and Result Presentation	Trippas	CIS manuscript, Ch. 3
3	45 min	Evaluating CIS Systems	Radlinski	CIS manuscript, Ch. 7
4	30 min	Understanding Conversational Interactions	Dalton	CIS manuscript, Ch. 4
5	30 min	Conversational Response Ranking and Generation	Dalton	CIS manuscript, Ch. 5
6	30 min	Mixed-Initiative Interactions	Zamani	CIS manuscript, Ch. 6
7	135 min	Hands-on with iCAST and OpenGRILL	Fischer, Owoicho & Rossetto	OpenGRILL, TaskMAD, and iCAST
8	15 min	Conclusion and Future Direction	Trippas	CIS manuscript, Ch. 8
All	6 hours			

Macaw [29] is an open-source platform for supporting conversational IR research. This Python-based platform extends multiple existing open-source search engines and question answering models. It provides user-friendly uni- and multi-modal interfaces. It is discussed in conjunction with the other tools.

Part 8: Conclusion and Future Direction. We conclude the tutorial by introducing open research problems and future directions in the area.

6 SUPPLEMENTAL MATERIAL

The tutorial website provides supplementary support materials including slides. This also includes specialized (simplified) software tools where appropriate.

ACKNOWLEDGMENTS

This work was supported in part by the Center for Intelligent Information Retrieval, in part by a 2021 Alexa Prize grant from Amazon, in part by a 2018 Google Faculty Research Award, and in part by an Engineering and Physical Sciences Research Council grant EP/V025708/1. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsors.

REFERENCES

- [1] Mohammad Aliannejadi and Johanne R. Trippas. 2022. Conversational Information Seeking: Theory and Evaluation: CHIIR 2022 Half Day Tutorial. In *ACM SIGIR Conference on Human Information Interaction and Retrieval* (Regensburg, Germany) (CHIIR '22). Association for Computing Machinery, New York, NY, USA, 365–366. <https://doi.org/10.1145/3498366.3505843>
- [2] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. Asking Clarifying Questions in Open-Domain Information-Seeking Conversations. In *Proceedings of the 42Nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR '19). ACM, New York, NY, USA, 475–484. <https://doi.org/10.1145/3331184.3331265>
- [3] Avishek Anand, Lawrence Cavendon, Hideo Joho, Mark Sanderson, and Benno Stein. 2020. Conversational Search (Dagstuhl Seminar 19461). Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [4] Jaime Arguello, Filip Radlinski, Hideo Joho, Damiano Spina, and Julia Kiseleva. 2018. Second International Workshop on Conversational Approaches to Information Retrieval (CAIR '18): Workshop at SIGIR 2018. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) (SIGIR '18). ACM, New York, NY, USA, 1435–1437. <https://doi.org/10.1145/3209978.3210199>
- [5] L. Azzopardi, M. Dubiel, M. Halvey, and J. Dalton. 2018. Conceptualizing agent-human interactions during the conversational search process. In *The Second International Workshop on Conversational Approaches to Information Retrieval* (CAIR '18).
- [6] Mikhail Burtsev, Aleksandr Chuklin, Julia Kiseleva, and Alexey Borisov. 2017. Search-Oriented Conversational AI (SCAI). In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval* (Amsterdam, The Netherlands) (ICTIR '17). ACM, New York, NY, USA, 333–334. <https://doi.org/10.1145/3121050.3121111>
- [7] Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wen-tau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. QuAC: Question Answering in Context. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Brussels, Belgium, 2174–2184. <https://doi.org/10.18653/v1/D18-1241>
- [8] Aleksandr Chuklin, Jeff Dalton, Julia Kiseleva, Alexey Borisov, and Mikhail Burtsev (Eds.). 2018. *Proceedings of the 2018 EMNLP Workshop SCAI: The 2nd International Workshop on Search-Oriented Conversational AI*. Association for Computational Linguistics, Brussels, Belgium. <https://www.aclweb.org/anthology/W18-5700>
- [9] J. Shane Culpepper, Fernando Diaz, and Mark D. Smucker. 2018. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *SIGIR Forum* 52, 1 (Aug. 2018), 34–90. <https://doi.org/10.1145/3274784.3274788>
- [10] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2019. TREC CAsT 2019: The Conversational Assistance Track Overview. (2019).
- [11] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2021. TREC CAsT 2021: The Conversational Assistance Track Overview. (2021).
- [12] Yashar Deldjoo, Johanne R. Trippas, and Hamed Zamani. 2021. Towards Multi-Modal Conversational Information Seeking. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, Canada) (SIGIR '21). ACM, New York, NY, USA.
- [13] Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2020. Tutorial on Conversational Recommendation Systems. In *Fourteenth ACM Conference on Recommender Systems*. 751–753.
- [14] Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2021. WSDM 2021 Tutorial on Conversational Recommendation Systems. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (Virtual Event, Israel) (WSDM '21). Association for Computing Machinery, New York, NY, USA, 1134–1136. <https://doi.org/10.1145/3437963.3441661>
- [15] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval* (Ann Arbor, MI, USA) (SIGIR '18). ACM, New York, NY, USA, 1371–1374. <https://doi.org/10.1145/3209978.3210183>
- [16] Jianfeng Gao, Chenyan Xiong, and Paul N. Bennett. 2020. Recent Advances in Conversational Information Retrieval. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (2020).
- [17] Wenqiang Lei, Xiangnan He, Maarten de Rijke, and Tat-Seng Chua. 2020. *Conversational Recommendation: Formulation, Methods, and Evaluation*. Association for Computing Machinery, New York, NY, USA, 2425–2428. <https://doi.org/10.1145/3397271.3401419>
- [18] Filip Radlinski, Krisztian Balog, Bill Byrne, and Karthik Krishnamoorthi. 2019. Coached Conversational Preference Elicitation: A Case Study in Understanding Movie Preferences. In *Proceedings of the Annual SIGDial Meeting on Discourse and Dialogue*.
- [19] Filip Radlinski and Nick Craswell. 2017. A Theoretical Framework for Conversational Search. In *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval* (Oslo, Norway) (CHIIR '17). Association for Computing Machinery, New York, NY, USA, 117–126. <https://doi.org/10.1145/3020165.3020183>
- [20] A. Ram, Rohit Prasad, C. Khatri, Anu Venkatesh, R. Gabriel, Q. Liu, J. Nunn, Behnam Hedayatnia, Ming Cheng, Ashish Nagar, E. King, Kate Bland, Amanda Wartick, Yi Pan, Han Song, Sk Jayadevan, Gene Hwang, and Art Pettigrew. 2018.

- Conversational AI: The Science Behind the Alexa Prize. *ArXiv abs/1801.03604* (2018).
- [21] Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A Conversational Question Answering Challenge. *Transactions of the Association for Computational Linguistics* 7 (2019), 249–266.
- [22] Alessandro Spezziorin, Jeffrey Dalton, and Anton Leuski. 2022. TaskMAD: A Platform for Multimodal Task-Centric Knowledge-Grounded Conversational Experimentation. *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (2022).
- [23] Damiano Spina, Jaime Arguello, Hideo Joho, Julia Kiseleva, and Filip Radlinski. 2019. CAIR'18: Second International Workshop on Conversational Approaches to Information Retrieval at SIGIR 2018. *SIGIR Forum* 52, 2 (Jan. 2019), 111–116. <https://doi.org/10.1145/3308774.3308793>
- [24] Paul Thomas, Daniel McDuff, Mary Czerwinski, and Nick Craswell. 2017. MISC: A data set of information-seeking conversations. In *Proceedings of the 1st International Workshop on Conversational Approaches to Information Retrieval*.
- [25] Johanne R. Trippas, Damiano Spina, Lawrence Cavedon, Hideo Joho, and Mark Sanderson. 2018. Informing the Design of Spoken Conversational Search: Perspective Paper. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval* (New Brunswick, NJ, USA) (*CHIIR '18*). Association for Computing Machinery, New York, NY, USA, 32–41. <https://doi.org/10.1145/3176349.3176387>
- [26] Johanne R. Trippas, Damiano Spina, Paul Thomas, Mark Sanderson, Hideo Joho, and Lawrence Cavedon. 2020. Towards a Model for Spoken Conversational Search. *Information Processing & Management* 57, 2 (2020), 102162.
- [27] Johanne R. Trippas and Paul Thomas. 2019. Data Sets for Spoken Conversational Search. In *Proceedings of the CHIIR 2019 Workshop on Barriers to Interactive IR Resources Re-use* (Glasgow, UK) (*BIIRRR@CHIIR '19*). 14–18.
- [28] Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, P. Bennett, Jianfeng Gao, and Zhiyuan Liu. 2020. Few-Shot Generative Conversational Query Rewriting. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (2020).
- [29] Hamed Zamani and Nick Craswell. 2020. Macaw: An Extensible Conversational Information Seeking Platform. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (*SIGIR '20*). Association for Computing Machinery, New York, NY, USA, 2193–2196. <https://doi.org/10.1145/3397271.3401415> <https://github.com/hamed-zamani/macaw>.
- [30] Hamed Zamani, Susan T. Dumais, Nick Craswell, Paul N. Bennett, and Gord Lueck. 2020. Generating Clarifying Questions for Information Retrieval. In *Proceedings of the 29th International Conference on World Wide Web* (Taipei, Taiwan) (*WWW '20*).
- [31] Hamed Zamani, Gord Lueck, Everest Chen, Rodolfo Quispe, Flint Luu, and Nick Craswell. 2020. MIMICS: A Large-Scale Data Collection for Search Clarification. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (Virtual Event, Ireland) (*CIKM '20*). Association for Computing Machinery, New York, NY, USA, 3189–3196. <https://doi.org/10.1145/3340531.3412772>
- [32] Hamed Zamani, Bhaskar Mitra, Everest Chen, Gord Lueck, Fernando Diaz, Paul N. Bennet, Nick Craswell, and Susan T. Dumais. 2020. Analyzing and Learning from User Interactions for Search Clarification. In *The 43rd International ACM SIGIR Conference on Research & Development in Information Retrieval* (Virtual Event, China) (*SIGIR '20*).
- [33] Hamed Zamani, Johanne R. Trippas, Jeff Dalton, and Filip Radlinski. 2022. Conversational Information Seeking. *Foundations and Trends in Information Retrieval* (2022). (to appear).
- [34] ChengXiang Zhai. 2020. *Interactive Information Retrieval: Models, Algorithms, and Evaluation*. Association for Computing Machinery, New York, NY, USA, 2444–2447. <https://doi.org/10.1145/3397271.3401424>
- [35] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. Towards Conversational Search and Recommendation: System Ask, User Respond. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management* (Torino, Italy) (*CIKM '18*). ACM, New York, NY, USA, 177–186. <https://doi.org/10.1145/3269206.3271776>