

Conversational Search and Recommendation: Introduction to the Special Issue

CLAUDIA HAUFF, Delft University of Technology, The Netherlands

JULIA KISELEVA, Microsoft, United States

MARK SANDERSON, RMIT University, Australia

HAMED ZAMANI*, University of Massachusetts Amherst, United States

YONGFENG ZHANG, Rutgers University, United States

Additional Key Words and Phrases: Conversational Search, Conversational Recommendation, Conversational Question Answering, Interactive Information Retrieval

ACM Reference Format:

Claudia Hauff, Julia Kiseleva, Mark Sanderson, Hamed Zamani, and Yongfeng Zhang. 2021. Conversational Search and Recommendation: Introduction to the Special Issue. *ACM Transactions on Information Systems* 1, 1, Article 1 (January 2021), 5 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

While conversational search and recommendation has roots in early Information Retrieval (IR) research, the recent advances in automatic voice recognition and conversational agents have created increasing interest in this area. This topic was recognized as an emerging research area in the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018) [Culpepper et al. 2018].¹ Conversational search and recommendation systems consist of multiple components, from user modeling to conversational understanding to query modeling to result presentation.

In recent years, the IR and related communities have witnessed a number of major contributions to the field of conversational search and recommendation. They include but are not limited to conversational search conceptualization (e.g., [Azzopardi et al. 2018; Deldjoo et al. 2021; Radlinski and Craswell 2017]), effective conversational query re-writing (e.g., [Yu et al. 2020]), generating and selecting clarifying questions (e.g., [Zamani et al. 2020a,c]), conversational preference elicitation (e.g., [Radlinski et al. 2019; Zhang et al. 2018]), and understanding user interactions with spoken conversational systems (e.g., [Trippas et al. 2018, 2020]). The growing body of work in this area has been supplemented by an increasing number of recent seminars (e.g., [Anand et al. 2020]), workshops (e.g., [Arguello et al. 2018; Burtsev et al. 2017; Chuklin et al. 2018; Spina et al. 2019]), tutorials (e.g., [Fu et al. 2020; Gao et al. 2018]), shared-tasks (e.g., [Dalton et al. 2019, 2020; Ram et al. 2018]), and datasets (e.g., [Aliannejadi et al. 2019; Choi et al. 2018; Radlinski et al.

*Corresponding author.

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

Authors' addresses: Claudia Hauff, c.hauff@tudelft.nl, Delft University of Technology, The Netherlands; Julia Kiseleva, julia.kiseleva@microsoft.com, Microsoft, United States; Mark Sanderson, mark.sanderson@rmit.edu.au, RMIT University, Australia; Hamed Zamani, zamani@cs.umass.edu, University of Massachusetts Amherst, United States; Yongfeng Zhang, yongfeng.zhang@rutgers.edu, Rutgers University, United States.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

Manuscript submitted to ACM

2019; Reddy et al. 2019; Thomas et al. 2017; Trippas and Thomas 2019; Zamani et al. 2020b]). Despite this great success, there are still numerous open problems in this domain that require in-depth investigation. Papers presented in this special issue target some of these important problems and suggest insightful future directions for the field.

2 OVERVIEW OF PAPERS

We received 27 submissions in response to the “call for papers” which was first announced in March 2020. This special issue presents the papers that were accepted for publication. It covers a wide range of topics related to conversational systems, including one paper on the theory of conversational interactions, three papers on conversational recommender systems, three papers on chatbots and dialogue systems, five papers on conversational search systems, two papers on evaluating conversational information seeking, one paper that presents a contrasting perspective of a semantic conversational system used with IoT devices, and one paper on diversifying responses in conversation models.

Thomas et al. [2021] in the first article of the issue provide a theoretical analysis of human-to-human conversations and connect the dots to recent advances in conversational search systems. In more detail, they examine well-established social communication theories within the context of conversations with humans in addition to agents and robots. Their work suggests initial guidelines and makes recommendations for future work on social communication, moves and structures, affect, style and alignment, simulations, data creation, and evaluation methodologies in the context of conversational IR systems.

In the second article of this issue, Li et al. [2021a] propose *Conversational Thompson Sampling*, a model for conversational recommendation systems that focuses on cold-start users. The work advances the state-of-the-art in conversational recommender systems. In the next article, Contreras et al. [2021] study conversational group recommender systems. with a web-based interface and perform extensive user studies to demonstrate the effectiveness of the proposed approach. A unique property of their approach is the exploitation of the interactions between users in a group during the recommendation sessions. Related to this research topic, Wilkinson et al. [2021] bridge the gap between conversational recommendation and trust and explainability. They focus on “why” and “why not” justifications for recommendations through conversational interactions in the domain of movie recommendation. This work contains interesting findings for conversational recommendation design. For example, the online experiments suggest that the studied “why” justifications (but not the “why not” justifications) have a significant impact on users’ perception of the recommendation engine.

Even though chatbots and dialogue systems are not mainly targeting information seeking tasks, such as search and recommendation, they are somewhat related to the broad area of conversational search and recommendation. Wei et al. [2021] design and develop an end-to-end neural chatbot architecture that is optimized to provide responses that are not only semantically relevant and coherent, but also take emotional information into account. This joint optimization leads to improvements in terms of both semantic- and emotional-level evaluation. On the other hand, Xu et al. [2021] focus on retrieval-based approaches to dialogue systems and propose a representation-interaction-matching framework. The intuition behind this work is to model the relation of neighboring elements and phrasal patterns in addition to long-range dependencies in dialogues. The paper presents strong performances on three public multi-turn dialogue response selection benchmarks. Related to this work, Li et al. [2021b] study personalized wording in dialogue history to improve retrieval-based dialogue systems.

Four of the accepted articles are related to conversational search systems. Musto et al. [2021] build a personal digital assistant, called MyrrorBot, that enables users to have personalized access to online services, such as music, videos, games, and news. A nice characteristic of MyrrorBot is that it enables users to query their own user profiles to

learn about the features encoded in their profile which improves the transparency of the personalization process. The Qualitative and quantitative evaluations of the model through a set of user studies demonstrate the effectiveness of the developed agent. Ren et al. [2021] study the concept of conversation with search engines, whose goal is to provide a natural language interface to the users that supports multi-turn interactions. They focus on addressing complex information needs that require multi-turn user-system interactions. They create and release the SaaC dataset (Search as a Conversation) and propose a pipeline for conversations with search engines. Lin et al. [2021] focus on conversational passage retrieval and propose two conversational query reformulation methods. The first method focuses on term importance estimation, while the other method rewrites the query using a sequence-to-sequence model. The outputs of these two methods are fused and led to 30% NDCG@3 improvement compared to the top performing submission at the TREC CAsT Track in 2019. Vakulenko et al. [2021] perform a large-scale analysis of mixed-initiative interactions in 15 publicly available datasets and uncover the limitations of existing benchmarks. Kiesel et al. [2021] conduct a through study of meta-information in conversational search. Exchange of meta-information in conversational search is necessary to complete the seeker's state of knowledge. The authors refine the existing definitions of meta-information for conversational search and conduct quantitative and qualitative analysis on MISC [Thomas et al. 2017] and SCS datasets [Trippas and Thomas 2019] for analyzing conversational interactions that contain meta-information.

Evaluating conversational search, and interactive IR in general, is challenging. Lipani et al. [2021] explore offline evaluation of conversational search systems. The work is based on the idea of subtopics which has been extensively used in evaluating search novelty and diversity. The proposed evaluation framework is based on a user model similar to the well-known geometric browsing model. Liu et al. [2021] evaluate the existing evaluation metrics used in the conversational search literature in terms of reliability, fidelity, and intuitiveness. This study shows that existing metrics correlates weakly with user satisfaction, and also suggests that existing session-based evaluation metrics can be used for measuring the quality of multi-turn conversational search with moderate concordance with user satisfaction.

Corno et al. [2021] study conversational interactions for the Internet of Things by exploring IF-THEN rules that can be deployed in different contexts. This paper presents a semantic conversational system, called HeyTAP, that takes into account the current user's intention, the connected entities owned by the user, and the user's long-term preferences.

In the last article included in this special issue, Yan et al. [2021] focus on improving the diversity of responses in conversation models. They achieve this goal by learning end-to-end dynamic representations for each input query conditioned on different response candidates. Evaluation on a wide range of human-to-human conversational data, including those collected from community question answering websites, forums, and micro-blog social networks, demonstrates the effectiveness of the proposed approach in terms of both appropriateness and diversity.

3 CONCLUSIONS

The papers included in this special issue cover a broad range of topics related to conversational search and recommendation, from theoretical analysis to algorithmic modeling to evaluation. Despite the growing body of work in conversational search and recommendation, it is important to note that most of the topics in this area are still under-researched. We are still yet to observe many breakthrough ideas in this domain in the coming years.

4 ACKNOWLEDGEMENTS

We thank all the authors who submitted their work to the special issue and all the reviewers who spent significant amount of time and effort to provide useful and thought-provoking feedback to the authors. We also thank Maarten de Rijke and Min Zhang, the last and the current Editor-in-Chiefs of the journal for their guidance.

REFERENCES

- 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 (SIGIR'19)*. ACM, New York, NY, USA, 475–484. <https://doi.org/10.1145/3331184.3331265>
- Avishke Anand, Lawrence Cavedon, Hideo Joho, Mark Sanderson, and Benno Stein. 2020. Conversational Search (Dagstuhl Seminar 19461). Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- 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 (SIGIR '18)*. ACM, New York, NY, USA, 1435–1437. <https://doi.org/10.1145/3209978.3210199>
- 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)*.
- 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 (ICTIR '17)*. ACM, New York, NY, USA, 333–334. <https://doi.org/10.1145/3121050.3121111>
- 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>
- 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>
- David Contreras, Maria Salamó, and Boratto Ludovico. 2021. Integrating Collaboration and Leadership in Conversational Group Recommender Systems. *ACM Transactions on Information Systems* (2021).
- Fulvio Corno, Luigi De Russis, and Alberto Monge Roffarello. 2021. From Users' Intentions to IF-THEN Rules in the Internet of Things. *ACM Transactions on Information Systems* (2021).
- 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>
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2019. TREC CAsT 2019: The Conversational Assistance Track Overview. (2019).
- Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2020. TREC CAsT 2020: The Conversational Assistance Track Overview. (2020).
- 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 (SIGIR '21)*. ACM, New York, NY, USA.
- Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2020. Tutorial on Conversational Recommendation Systems. In *Fourteenth ACM Conference on Recommender Systems*. 751–753.
- 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 (SIGIR '18)*. ACM, New York, NY, USA, 1371–1374. <https://doi.org/10.1145/3209978.3210183>
- Johannes Kiesel, Lars Meyer, Martin Potthast, and Benno Stein. 2021. Meta-Information in Conversational Search. *ACM Transactions on Information Systems* (2021).
- Juntao Li, Chang Liu, Chongyang Tao, Zhangming Chan, Dongyan Zhao, Min Zhang, and Rui Yan. 2021b. Dialogue History Matters! Personalized Response Selection in Multi-turn Retrieval-based Chatbots. *ACM Transactions on Information Systems* (2021).
- Shijun Li, Wenqiang Lei, Qingyun Wu, Xiangnan He, Peng Jiang, and Tat-Seng Chua. 2021a. Seamlessly Unifying Attributes and Items: Conversational Recommendation for Cold-Start Users. *ACM Transactions on Information Systems* (2021).
- Sheng-Chieh Lin, Jheng-Hong Yang, Rodrigo Nogueira, Ming-Feng Tsai, Chuan-Ju Wang, and Jimmy Lin. 2021. Multi-Stage Conversational Passage Retrieval: An Approach to Fusing Term Importance Estimation and Neural Query Rewriting. *ACM Transactions on Information Systems* (2021).
- Aldo Lipani, Ben Carterette, and Emine Yilmaz. 2021. How Am I Doing?: Evaluating Conversational Search Systems Offline. *ACM Transactions on Information Systems* (2021).
- Zeyang Liu, Ke Zhou, and Max L. Wilson. 2021. Meta-Evaluation of Conversational Search Evaluation Metrics. *ACM Transactions on Information Systems* (2021).
- Cataldo Musto, Fedelucio Narducci, Marco Polignano, Marco De Gemmis, Pasquale Lops, and Giovanni Semeraro. 2021. MyrrorBot: a Digital Assistant Based on Holistic User Models for Personalized Access to Online Services. *ACM Transactions on Information Systems* (2021).
- 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*.
- 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 (CHIIR '17)*. Association for Computing Machinery, New York, NY, USA, 117–126. <https://doi.org/10.1145/3020165.3020183>
- 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).

- 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.
- Pengjie Ren, Zhumin Chen, Zhaochun Ren, Evangelos Kanoulas, Christof Monz, and Maarten de Rijke. 2021. Conversations with Search Engines: SERP-based Conversational Response Generation. *ACM Transactions on Information Systems* (2021).
- 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>
- Paul Thomas, Mary Czerwinski, Daniel McDuff, and Nick Craswell. 2021. Theories of conversation for conversational IR. *ACM Transactions on Information Systems* (2021).
- 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*.
- 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 (CHIIR '18)*. Association for Computing Machinery, New York, NY, USA, 32–41. <https://doi.org/10.1145/3176349.3176387>
- 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.
- 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 (BIIRRR@CHIIR '19)*. 14–18.
- Svitlana Vakulenko, Evangelos Kanoulas, and Maarten de Rijke. 2021. A Large-Scale Analysis of Mixed Initiative in Information-Seeking Dialogues for Conversational Search. *ACM Transactions on Information Systems* (2021).
- Wei Wei, Jiayi Liu, Xianling Mao, Guibing Guo, Feida Zhu, Pan Zhou, Yuchong Hu, and Shanshan Feng. 2021. Target Guided Emotion Aware Chat Machine. *ACM Transactions on Information Systems* (2021).
- Darcia Wilkinson, Öznur Alkan, Q. Vera Liao, Massimiliano Mattetti, Inge Vejsbjerg, Bart P. Knijnenburg, and Elizabeth Daly. 2021. Why or Why Not? The Effect of Justification Styles on Chatbot Recommendations. *ACM Transactions on Information Systems* (2021).
- Ruijian Xu, Chongyang Tao, Jiazhan Feng, Wei Wu, Rui Yan, and Dongyan Zhao. 2021. Response Ranking with Multi-Types of Deep Interactive Representations in Retrieval-based Dialogues. *ACM Transactions on Information Systems* (2021).
- Rui Yan, Weiheng Liao, Dongyan Zhao, and Ji-Rong Wen. 2021. Multi-Response Awareness for Retrieval-Based Conversations: Respond with Diversity via Dynamic Representation Learning. *ACM Transactions on Information Systems* (2021).
- 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).
- Hamed Zamani, Susan T. Dumais, Nick Craswell, Paul N. Bennett, and Gord Lueck. 2020a. Generating Clarifying Questions for Information Retrieval. In *Proceedings of the 29th International Conference on World Wide Web (WWW '20)*.
- Hamed Zamani, Gord Lueck, Everest Chen, Rodolfo Quispe, Flint Luu, and Nick Craswell. 2020b. MIMICS: A Large-Scale Data Collection for Search Clarification. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 3189–3196. <https://doi.org/10.1145/3340531.3412772>
- Hamed Zamani, Bhaskar Mitra, Everest Chen, Gord Lueck, Fernando Diaz, Paul N. Bennet, Nick Craswell, and Susan T. Dumais. 2020c. Analyzing and Learning from User Interactions for Search Clarification. In *The 43rd International ACM SIGIR Conference on Research & Development in Information Retrieval (SIGIR '20)*.
- 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 (CIKM '18)*. ACM, New York, NY, USA, 177–186. <https://doi.org/10.1145/3269206.3271776>