

CAIR'18: Second International Workshop on Conversational Approaches to Information Retrieval at SIGIR 2018

Damiano Spina
RMIT University
damiano.spina@rmit.edu.au

Jaime Arguello
University of North Carolina at Chapel Hill
jarguell@email.unc.edu

Hideo Joho
University of Tsukuba
hideo@slis.tsukuba.ac.jp

Julia Kiseleva
University of Amsterdam
j.kiseleva@uva.nl

Filip Radlinski
Google
filiprad@google.com

Abstract

The Second International Workshop on Conversational Approaches to Information Retrieval (CAIR'18) was held on July 12th, 2018 in Ann Arbor, Michigan, USA in association with SIGIR 2018, the 41st International ACM SIGIR Conference on Research and Development in Information Retrieval. CAIR'18 brought together academic and industry researchers to present and discuss new research on conversational approaches to search. This report provides a summary of the workshop consisting of two keynote talks, five oral presentations, and a joint panel discussion with the KG4IR workshop.

1 Introduction

The Second International Workshop on Conversational Approaches to Information Retrieval (CAIR'18)¹ [2] was held in Ann Arbor, Michigan, USA on July 12th, 2018 as a full-day workshop associated with SIGIR 2018. As a follow-up to the first edition [8], the workshop was organized in response to the growth in speech interfaces, particularly on mobile devices, and to Personal Digital Assistants (PDAs) such as Apple Siri, Microsoft Cortana, Google Assistant, Amazon Alexa and others, and the opportunities afforded by this burgeoning mode of interaction.

The Call for Papers specifically welcomed research on techniques that support complex and multi-turn user-machine dialogues for information access and retrieval, and multi-modal interfaces for interacting with such systems. It also welcomed work addressing all modalities of conversation, including speech-based, text-based, and multi-modal interaction, as well as studies of human-human interaction (e.g., collaborative search) that can inform the design of conversational search applications, and work on evaluation of conversational approaches.

¹<https://sites.google.com/view/cair-ws>

2 Program at a Glance

The workshop program included two keynotes, five oral presentations in two sessions, poster presentations, and a panel discussion, jointly organized with the workshop on Knowledge Graphs and Semantics for Text Retrieval, Analysis, and Understanding (KG4IR)².

The structure of the CAIR'18 workshop was as follows.

- Keynote Talk I: *Steps Toward Semantic Search* (Brian Strope, Google)
- Oral Session 1
- Oral Session 2
- Keynote Talk II: *Conversational Search with Voice-based Assistants: The case of News Search Enrico uh Suggestion* (Eugene Agichtein, Emory University)
- Poster Discussions
- Joint Panel together with the KG4IR Workshop + Closing

3 Keynote Speakers

Two keynote presenters were invited, one from industry and one from academia.

The first keynote speaker **Brian Strope**³ is a research scientist and manager for Ray Kurzweil's research efforts at Google. Earlier at Google, Brian helped build Google's speech recognition technology, and he worked for several years as a research engineer at Nuance Communications. His PhD dissertation at UCLA was in auditory modeling, and he studied signal processing and music at Brown. In his talk, Brian presented approaches to extract semantic structures from large amounts of unstructured data at scale. Techniques presented in this talk include a large-scale semantic similarity measurement of utterances by examining whether utterances had very similar responses in online forum conversations. Brian showed the effectiveness of these techniques in the context of reply suggestion to e-mails.

The second keynote speaker **Eugene Agichtein**⁴ is an Associate Professor of Computer Science at Emory University, where he founded and leads the Intelligent Information Access Laboratory (IR Lab). Eugene's research spans the areas of information retrieval, natural language processing, data mining, and human computer interaction. In his talk, Eugene presented the lessons learned from participating in the Amazon Alexa Prize, a global university competition for conversational AI.

4 Papers: Main Themes

The workshop program had a total of five oral presentations in two sessions. The findings presented in those papers can be grouped into three main themes: Support, User Behavior, and Techniques.

²<https://kg4ir.github.io/>

³<https://goo.gl/5WEh4i>

⁴<http://www.mathcs.emory.edu/~eugene/>

4.1 Support: What should a Conversational IR system be able to do?

An important question for conversational IR is: What should the system be able to do? In a position paper, Azzopardi et al. [5] describe different requirements of a fully conversational IR system. These requirements were derived by enumerating the different types of user and agent *actions* associated with conversational search. For example, user actions include: (1) revealing information about the current information need (CIN), (2) modifying the CIN (e.g., broadening/narrowing the scope), (3) navigating the information found by the agent, (4) asking what the agent knows/understands about the CIN, and (5) closing a search session. In contrast, agent actions include: (1) eliciting information, (2) asking for clarification, (3) presenting information found (e.g., lists, summaries, comparisons), (4) suggesting alternative information needs (AINs), and (5) explaining actions. Based on these different user/agent actions, the paper proposes different predictive tasks that could be potentially investigated separately, including: (1) predicting what agent action to do next (inquire, reveal, traverse, suggest), (2) extracting information from a user’s turn to update the CIN representation, (3) deciding which question(s) to ask to elicit information, (4) predicting when to ask for clarification (versus *infer* unknown information about the user’s goals), and (5) predicting when to suggest an AIN based on the current session.

4.2 User Behavior: How do users interact with a Conversational IR system?

Several papers presented at the workshop reported on user studies that investigated different research questions related to conversational approaches to IR.

Avula [3] elaborates on the benefits and challenges of implementing a “Wizard of Oz” (WoZ) methodology to study conversational IR systems. A WoZ study is one in which participants interact with one or more systems they believe to be autonomous, but are actually operated (partially or completely) by an unseen human. In conversational IR, WoZ studies enable researchers to study systems that are not yet technologically feasible and/or may be prone to errors that are unrelated to the research questions being examined. In this paper, Avula reflects on lessons learned during a previously published WoZ study [4] that investigated two different ways for chatbots to intervene in Slack conversations to provide assistance. Avula reflects on several decisions made during that study that proved to be useful. These included: (1) developing the infrastructure for the “Wizard” to make decisions and respond quickly, (2) giving the “Wizard” clear rules on how to interact with participants in order to ensure consistency within experimental conditions, and (3) giving participants clear instructions about the system’s capabilities to prevent participants in the same condition from having different expectations (and therefore different experiences).

Dubiel et al. [7] described a “Wizard of Oz” study that compared a slot-filling system against a fully conversational system. In this study, participants completed four “flight booking” tasks (two with each system). During each task, participants had to book a flight under different constraints. The slot-filling system interacted with participants using a scripted dialogue, while the fully conversational system interacted with participants using human-like, spontaneous dialogue. In the fully conversational condition, participants completed tasks using fewer (and shorter) utterances (due to less repetition), achieved better outcomes (e.g., found cheaper flights), reported lower levels of workload, reported higher levels of system

usability, and produced utterances with more positive sentiment.

Kaushik and Jones [9] described two studies that analyzed participants' search activities during complex tasks, and related these two different processes associated with Vakkari's *search as learning* model [10]: restructuring knowledge, stabilizing knowledge structures, and assimilating new information with existing structures. Participants exhibited different behaviors while pursuing different learning objectives. The authors propose that "search as learning" may be a useful framework/perspective from which to design and evaluate conversational IR systems. Future work is needed to understand how a conversational IR system may support different learning goals.

4.3 Techniques: How can the functionality of a Conversational IR system be advanced?

Wambua et al. [11] presented a conversational IR system in the healthcare domain—a system developed to support patients and caregivers with information needs related to dementia and Alzheimer's disease. The authors developed a system that allows users to iteratively refine their searches by answer questions about different facet-value pairs associated with the collection. Facet-value pairs were developed in advance and predicted using supervised machine learning (e.g., audience={patients, caregivers, healthcare professionals, researchers}). Similar to the ID3 decision tree algorithm, after an initial query, the system iteratively asks the user about the facet associated with the greatest *information gain* given previous responses.

5 Joint Panel together with the KG4IR Workshop

Given the overlap of research topics between the CAIR and the KG4IR workshops, the organizers of both workshop decided to have a joint panel discussion.

The panel brought together researchers that provided strong expertise on topics related to both workshops, including: knowledge graphs, intelligent assistants, question answering, knowledge processing, entity linking, conversational search, recommendation, and speech recognition.

Chairs

- Laura Dietz (University of New Hampshire)
- Jaime Arguello (University of North Carolina at Chapel Hill)

Panelists

- Eugene Agichtein (Emory University)
 - Soumen Chakrabarti (IIT Bombay)
 - Jeffrey Dalton (University of Glasgow)
 - Brian Strope (Google)
 - Scott Wen-tau Yih (Allen Institute for Artificial Intelligence)
-

Summary. The panel started discussing the different advantages and disadvantages of using semantic technologies and knowledge graphs in conversational systems. For instance, knowledge graphs may provide an effective way to keep track of the state in conversational systems. One challenge of using semantic technologies in conversational scenarios is that knowledge bases are typically incomplete, objective and context-unaware.

The lack of publicly available resources for Conversational IR was one of the main topics of discussions. On one hand, industry has access to user interactions with existing systems that are not easily releasable to researchers in academia. Here, initiatives such as the Amazon Alexa Prize seems to be a good approach. On the other hand, industry and academia face a common challenge which is “there is no data available for systems have not yet been implemented”. For instance, problems such as *monetization* (e.g., present sponsored advertisements via a speech-only channel), *internationalization* or *personalization* of Conversational IR systems remain as open challenges.

6 Conclusions

The active participation of the community – from both industry and academia – at the Second International Workshop on Conversational Approaches to Information Retrieval (CAIR’18) demonstrates the growing interest to the problem of conversational IR, which is far from being a solved problem. In fact, progress in this area has been to some extent slower than predicted at SWIRL 2012 [1]. Nevertheless, the recent SWIRL 2018 [6] has reiterated Conversational IR as one of the main challenges for the next years. We envisage that the creation of new resources and the adaptation of new experimental methodologies will accelerate the advancement of Conversational IR.

The CAIR’18 workshop organizers would like to acknowledge the following people for their contributions to the successful workshop: Brian Strope (Google), Eugene Agichtein (Emory University), Jeffrey Dalton (University of Glasgow), Laura Dietz (University of New Hampshire), Scott Wen-tau Yih (Allen Institute for Artificial Intelligence), Soumen Chakrabarti (IIT Bombay), Lawrence Cavedon (RMIT University), Fernando Diaz (Microsoft), Dilek Hakkani-Tur (Google), Mark Sanderson (RMIT University), Milad Shokouhi (Microsoft), Ido Guy (eBay Research), Claudia Hauff (TU Delft), Gareth Jones (Dublin City University), Ben Lambert (Spotify), Jiepu Jiang (Virginia Tech), Karthik Raghunathan (MINDMELD), Paul Thomas (Microsoft), Johanne R. Trippas (RMIT University), and Imed Zitouni (Microsoft Research).

References

- [1] James Allan, Bruce Croft, Alistair Moffat, and Mark Sanderson. Frontiers, Challenges, and Opportunities for Information Retrieval: Report from SWIRL 2012 the Second Strategic Workshop on Information Retrieval in Lorne. *SIGIR Forum*, 46(1):2–32, 2012. ISSN 0163-5840. doi: 10.1145/2215676.2215678.
 - [2] Jaime Arguello, Filip Radlinski, Hideo Joho, Damiano Spina, and Julia Kiseleva. Second International Workshop on Conversational Approaches to Information Retrieval (CAIR’18): Workshop at SIGIR 2018. In *The 41st International ACM SIGIR Confer-*
-

-
- ence on Research & Development in Information Retrieval, SIGIR '18, pages 1435–1437, 2018. doi: 10.1145/3209978.3210199.
- [3] Sandeep Avula. Wizard of Oz: Protocols and Challenges in Studying Searchbots to Support Collaborative Search. In *SIGIR 2nd International Workshop on Conversational Approaches to Information Retrieval (CAIR'18)*, 2018.
- [4] Sandeep Avula, Gordon Chadwick, Jaime Arguello, and Robert Capra. SearchBots: User Engagement with ChatBots During Collaborative Search. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval, CHIIR '18*, pages 52–61. ACM, 2018.
- [5] Leif Azzopardi, Mateusz Dubiel, Martin Halvey, and Jeffery Dalton. Conceptualizing Agent-Human Interactions During the Conversational Search Process. In *SIGIR 2nd International Workshop on Conversational Approaches to Information Retrieval (CAIR'18)*, 2018.
- [6] J. Shane Culpepper, Fernando Diaz, and Mark D. Smucker. Research Frontiers in Information Retrieval: Report from the Third Strategic Workshop on Information Retrieval in Lorne (SWIRL 2018). *SIGIR Forum*, 52(1):34–90, 2018. ISSN 0163-5840. doi: 10.1145/3274784.3274788.
- [7] Mateusz Dubiel, Martin Halvey, Leif Azzopardi, and Sylvain Darronat. Investigating How Conversational Search Agents Affect Users Behaviour, Performance and Search Experience. In *SIGIR 2nd International Workshop on Conversational Approaches to Information Retrieval (CAIR'18)*, 2018.
- [8] Hideo Joho, Lawrence Cavedon, Jaime Arguello, Milad Shokouhi, and Filip Radlinski. CAIR'17: First International Workshop on Conversational Approaches to Information Retrieval at SIGIR 2017. *SIGIR Forum*, 51(3):114–121, 2018. ISSN 0163-5840. doi: 10.1145/3190580.3190598.
- [9] Abhishek Kaushik and Gareth Jones. Exploring Current User Web Search Behaviours in Analysis Tasks to be Supported in Conversational Search. In *SIGIR 2nd International Workshop on Conversational Approaches to Information Retrieval (CAIR'18)*, 2018.
- [10] Pertti Vakkari. Searching as Learning: A Systematization Based on Literature. *Journal of Information Science*, 42(1):7–18, 2016.
- [11] Muuo Wambua, Stefania Raimondo, Jennifer Boger, Jan Polgar, Hamidreza Chinaei, and Frank Rudzicz. Interactive Search through Iterative Refinement. In *SIGIR 2nd International Workshop on Conversational Approaches to Information Retrieval (CAIR'18)*, 2018.
-