

Report on Supporting and Understanding of Conversational Dialogues workshop (SUD 2021) at WSDM 2021

Debasis Ganguly
University of Glasgow, UK
Debasis.Ganguly@glasgow.ac.uk

Gareth J.F. Jones
Dublin City University, Ireland
gareth.jones@dcu.ie

Procheta Sen
University of Glasgow, UK
Procheta.Sen@glasgow.ac.uk

Manisha Verma
Yahoo! Research, VerizonMedia, USA
manishav@verizonmedia.com

Dipasree Pal
Indian Statistical Institute, India
dipasree.pal@gmail.com

Abstract

This report describes the workshop on Supporting and Understanding of (multi-party) conversational Dialogues (SUD) organized as a part of the Web Search and Data Mining conference (WSDM) 2021. The aim of SUD workshop was to encourage researchers to investigate automated methods to analyze and understand conversations. We also discuss the release of a dataset that would be useful in IR research on conversations. The dataset was constructed to support the data challenge in SUD workshop and its precursor event - the Retrieval from Conversational Dialogues (RCD) track at the Forum of Information Retrieval and Evaluation (FIRE) 2020.

1 Introduction and Motivation for the Workshop

Conversational information retrieval (IR) is an upcoming area of research fueled by the rise in daily usage of smart assistants such as Alexa or Cortana [Hashemi et al., 2018; Ahmadvand et al., 2018]. While, traditional IR involves retrieving documents relevant to a user query, conversational IR, on the other hand, is useful to retrieve relevant documents or passages that address questions arising during interactive chat sessions. An example use-case is an automated customer-care chatbot that seeks to answer a specific user query (e.g. ‘How do I restart my Macbook?’) responding with a relevant answer, which could be formulated by either directly retrieving previous responses or constructing them automatically from retrieving candidate passages.

It has been established that IR systems can be used to retrieve relevant information either to generate replies in conversations [González-Garduño et al., 2019], or to add more context about a particular topic during an interactive dialogue between two entities [Dubiel et al., 2020]. The aim of the Supporting and Understanding of Dialogues (SUD) workshop [Ganguly et al., 2021] at the Web Search and Data Mining conference (WSDM’21) workshop was to encourage researchers to submit work that investigates various methods to analyze, understand or proactively support conversations.

Different from the existing notion of conversational IR that typically involves a human user with information needs and an automated agent seeking to find relevant answers to these information needs, the topic that particularly motivated us to conduct the SUD workshop was to explore ways of leveraging search systems to better facilitate the comprehension of conversational exchanges between *two or more humans*. As a precursor to the SUD workshop, we organized the Retrieval from Conversational Dialogues (RCD) track at the Forum of Information Retrieval and Evaluation (FIRE) 2020, where the task for the participants was to retrieve a set of relevant Wikipedia passages that may be useful for a better comprehension of the topics of a dialogue.

The SUD workshop had two tracks - a regular research track and a data challenge track which invited research on contextualizing movie dialogues using the dataset that was released during the RCD workshop [Ganguly et al., 2020]. The workshop program included several talks. There were two keynote talks - one from the academia and the other from the industry, a research paper presentation, and two invited talks. The workshop concluded with a panel discussion on the future directions of research on conversational IR.

2 Dataset Released for Conversational IR

As a part of the RCD track and the SUD workshop, we compiled a dataset comprised of selected excerpts from movie dialogues as queries (the dialogue selection process ensuring that the topics involve elements of general knowledge likely to be found in Wikipedia), and a set of judged relevant Wikipedia passages (obtained with depth-20 pooling).

If you let me, I can be your greatest ally. Take the **acacia tree**... in East Africa. It is the most... I’m not interested. Allow me to dose. As a sign of good faith we wish... The suitcase isn’t filled with fifties or gold or diamonds. Just silicon. A **Ming Mecca chip**.

Figure 1: A sample excerpt of dialogue from the movie *Pi* annotated with *topics* (spans of text) that potentially require *contextualization* for better comprehension of the dialogue.

Figure 1 shows an example conversation between two persons taken from the script of the movie ‘Pi’¹). The objective of an automated conversational assistance agent would be to first identify the concepts or entities that are potential candidates requiring further contextualization for instance the ones such as *acacia tree* and *Ming Mecca chip* (shown highlighted in the figure). Following an identification of the concepts, a search system can then formulate queries and retrieve potentially relevant information in the form of passages from Wikipedia, e.g. about the natural

¹<https://www.imdb.com/title/tt0138704/>

habitat and benefits of an acacia tree, or a passage defining what a Ming Mecca chip actually is – a modular synthesizer to manipulate retro video games.

The dataset² is freely available for research purposes. The target collection is available in two forms - as a Lucene index and as a tab separated text file containing a mapping of the unique ids and the passages. These passage ids are the ones that are used to identify the relevant passages in the qrels file. The queries are contained in a TREC formatted topic file comprised of 49 queries in total. The title field of each topic contains the annotated text span, whereas the description field contains the entire dialog excerpt.

3 Keynote Talks

There were two keynote talks in the workshop, one from the academia and the other from the industry. Details follow.

3.1 Keynote 1

We are grateful to Jeff Dalton of the University of Glasgow, who delivered the first keynote, mainly sharing his experience of organizing the ‘TREC Conversational Assistance Track (CAST)’ [Dietz et al., 2018], a conversational search benchmark at TREC³. More details follow.

Title: Conversation Search: Current Challenges and Future Directions.

Speaker Biography: Dr Jeff Dalton is an Assistant Professor (Lecturer) in the School of Computing Science at the University of Glasgow where he leads the Glasgow Representation and Information Learning Lab (GRILL). His research focuses on text understanding and conversational information seeking. He completed his PhD at the University of Massachusetts Amherst in the Center for Intelligent Information Retrieval. Later in Google Research, he worked on Information Extraction as part of the Knowledge Discovery Team (Knowledge Vault) and language understanding in the Assistant Response Ranking team. He is the lead organizer for the TREC Conversational Assistance Track (CAST)⁴ and previously helped organize the Complex Answer Retrieval track. He is the recipient of a prestigious UKRI Turing AI Acceleration Fellowship and received research awards from Google, Amazon, and Bloomberg. He holds multiple patents in retrieval and question answering systems.

Abstract: The talk examined some of the fundamental challenges in Conversational Information Seeking (CIS). It described lessons learned developing recent conversational search benchmarks (TREC CAST and ConvAI3) and open research challenges. It also presents work from the Glasgow Representation and Information Learning Lab (GRILL) on neural models for conversational rewriting, entity-based ranking, and the use of feedback and initiative in CIS.

²<https://rcd2020firetask.github.io/RCD2020FIRETASK/#dataset>

³<https://trec.nist.gov/>

⁴<http://treccast.ai>

3.2 Keynote 2

The second keynote was delivered by Surbhi Rathore, CEO of a startup organization that intends to develop AI-driven tools for management of conversational content. More details about the speaker and the talk abstract follow.

Title: Conversation intelligence for the growing digital ecosystem.

Speaker Biography: Surbhi is the CEO and Co-founder of Symbbl.ai⁵. Symbbl.ai is bringing to life her vision for a programmable platform that empowers developers and businesses to build unique conversational experiences without the hassle of upfront training or building their in-house data science expertise. She co-founded Symbbl.ai almost 2 years ago and backed by Amazon⁶. Symbbl.ai raised an early-stage venture round of \$6.5M and deployed capital to grow a team of tech enthusiasts to 30 people primarily distributed between India and Seattle. She comes with experience from technical and customer-obsessed roles in both startups and enterprises such as Nevis Networks⁷ and Amdocs⁸. Before co-founding Symbbl.ai, she worked in the Conversational AI space with a focus on delivering value to Telecommunication users. She is an advocate for Women in AI with a personal mission to inspire more women to work in Data Science. In her free time, she loves to travel and spend time with her remote and distributed family.

Abstract: Now, more than ever, when the whole world is communicating on digital channels, the dangers of disconnected data silos and lost knowledge are more than ever. The talk was about applying AI/ML to human conversations at scale so that builders can strategize product strategy early in the life cycle of their business and future-proof growth. The session also entails several dimensions of conversation intelligence that can amplify customer experiences, super-power productivity and drive growth opportunities for businesses for several mission critical use cases.

4 Other Talks

In addition to the two keynotes, there was one talk from the authors of the only accepted paper in the regular research paper track. This talk was followed by two more invited talks, the first by Emine Yilmaz of University College, London (UCL), and the second by Pabitra Mitra of the Indian Institute of Technology, Kharagpur (IIT-KGP). Details follow.

4.1 Accepted Paper

Title: Explaining Outcomes of Multi-Party Dialogues using Causal Learning.

Authors: *Priyanka Sinha, Pabitra Mitra, Antonio Anastasio Bruto da Costa and Nikolaos Kekatos*

⁵<https://symbbl.ai/>

⁶<https://www.amazon.com/>

⁷<https://nevisnetworks.com/>

⁸<https://www.amdocs.com/>

Abstract: Multi-party dialogues are common in enterprise social media on technical as well as non-technical topics. Outcome of a conversation may be positive or negative. It is important to analyse why a dialogue ends with a particular sentiment from the point of view of conflict analysis as well as future collaboration design. We propose an explainable time series mining algorithm for such analysis. A dialogue is represented as an attributed time series of occurrences of keywords, EMPATH [Fast et al., 2016] categories, and inferred sentiments at various points in its progress. A decision tree based classifier is used for predicting the outcome sentiment. Interpretable rules mined from the classifier are used to explain the prediction. Experimental results are presented for the enterprise social media posts in a large company. [Sinha et al., 2021]

4.2 Invited Talks

Emine Yilmaz of UCL summarized the role of modeling tasks for conversational AI systems in general, and conversational information seeking, in particular. Her talk covered some of her recent work in modeling tasks across user search sessions, e.g., using a non-parametric approach to identify the different tasks of a user across search sessions [Mehrotra and Yilmaz, 2017a,b], and also shared her experience of organizing the TREC task on modeling tasks from search sessions [Kanoulas et al., 2017].

Following Emine’s talk, Pabitra Mitra of IIT-KGP stressed the importance of developing conversation-based user interfaces as a primary form of interaction with AI tools supporting agriculture. This is because in a developing nation, such as India, this would allow farmers, a majority of whom lack adequate formal education, to leverage from the potential benefits of the AI tools for better crop management [Jain et al., 2018].

5 Panel Discussion on Future of Conversational IR

The workshop was concluded with a short (15 minutes) panel discussion, involving the organizers, the invited speakers and the participants. The main topics discussed were the following.

1. What are the meaningful tasks and the necessary datasets that we could envision to develop for carrying out further research in conversational systems?
2. What can be the different future use cases of conversational systems?
3. What is the role of modeling a user’s information seeking behaviour in supporting conversations?
4. What are the different aspects of a conversational system that an early researcher introduced to the subject (e.g., a post-graduate or a PhD. student) should focus on?

Acknowledgements

The organizers express their gratitude to the keynote speakers and the invited speakers for their invaluable contributions to the workshop.

References

- Ali Ahmadvand, Ingyu Choi, Harshita Sahijwani, Justus Schmidt, Mingyang Sun, Sergey Volokhin, Zihao Wang, and Eugene Agichtein. Emory irisbot: An open-domain conversational bot for personalized information access. *Alexa Prize Proceedings*, 2018.
- Laura Dietz, Ben Gamari, Jeff Dalton, and Nick Craswell. TREC complex answer retrieval overview. In Ellen M. Voorhees and Angela Ellis, editors, *Proceedings of the Twenty-Seventh Text REtrieval Conference, TREC 2018, Gaithersburg, Maryland, USA, November 14-16, 2018*, volume 500-331 of *NIST Special Publication*. National Institute of Standards and Technology (NIST), 2018.
- M. Dubiel, Martin Halvey, L. Azzopardi, D. Anderson, and Sylvain Daronnat. Conversational strategies : impact on search performance in a goal-oriented task. In *The Third International Workshop on Conversational Approaches to Information Retrieval*, 2020.
- Ethan Fast, Binbin Chen, and Michael S. Bernstein. Empath: Understanding topic signals in large-scale text. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, page 4647–4657. Association for Computing Machinery, 2016.
- Debasis Ganguly, Dipasree Pal, Manisha Verma, and Procheta Sen. Overview of rcd-2020, the fire-2020 track on retrieval from conversational dialogues. In *Proceedings of Forum for Information Retrieval Evaluation, 2020*, page 33–36, 2020.
- Debasis Ganguly, Manisha Verma, Procheta Sen, Dipasree Pal, and Gareth J.F. Jones. Overview of the Supporting and Understanding of Conversational Dialogues (SUD) Workshop. In *Proceedings of Web Search And Data Mining, 2021*, page 1163–1164, 2021.
- Ana Valeria González-Garduño, Isabelle Augenstein, and Anders Søgaard. Retrieval-based goal-oriented dialogue generation. *CoRR*, abs/1909.13717, 2019. URL <http://arxiv.org/abs/1909.13717>.
- Seyyed Hadi Hashemi, Kyle Williams, Ahmed El Kholy, Imed Zitouni, and Paul A Crook. Measuring user satisfaction on smart speaker intelligent assistants using intent sensitive query embeddings. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1183–1192, 2018.
- Mohit Jain, Pratyush Kumar, Ishita Bhansali, Q. Vera Liao, Khai N. Truong, and Shwetak N. Patel. Farmchat: A conversational agent to answer farmer queries. *Proceedings of ACM Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(4):170:1–170:22, 2018.
- Evangelos Kanoulas, Emine Yilmaz, Rishabh Mehrotra, Ben Carterette, Nick Craswell, and Peter Bailey. TREC 2017 tasks track overview. In Ellen M. Voorhees and Angela Ellis, editors, *Proceedings of The Twenty-Sixth Text REtrieval Conference, TREC 2017, Gaithersburg, Maryland, USA, November 15-17, 2017*, volume 500-324 of *NIST Special Publication*. National Institute of Standards and Technology (NIST), 2017.

-
- Rishabh Mehrotra and Emine Yilmaz. Extracting hierarchies of search tasks & subtasks via a bayesian nonparametric approach. In Noriko Kando, Tetsuya Sakai, Hideo Joho, Hang Li, Arjen P. de Vries, and Ryen W. White, editors, *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, Shinjuku, Tokyo, Japan, August 7-11, 2017*, pages 285–294. ACM, 2017a.
- Rishabh Mehrotra and Emine Yilmaz. Task embeddings: Learning query embeddings using task context. In Ee-Peng Lim, Marianne Winslett, Mark Sanderson, Ada Wai-Chee Fu, Jimeng Sun, J. Shane Culpepper, Eric Lo, Joyce C. Ho, Debora Donato, Rakesh Agrawal, Yu Zheng, Carlos Castillo, Aixin Sun, Vincent S. Tseng, and Chenliang Li, editors, *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 - 10, 2017*, pages 2199–2202. ACM, 2017b.
- Priyanka Sinha, Pabitra Mitra, Antonio Anastasio Bruto da Costa, and Nikolaos Kekatos. Explaining outcomes of multi-party dialogues using causal learning. In *Workshop on Supporting and Understanding of Conversational Dialogues (SUD-2021) at WSDM 2021*, 2021. URL <https://arxiv.org/abs/2105.00944>.