

WSDM 2021 Tutorial on Conversational Recommendation Systems

Zuohui Fu
Dept of Computer Science
Rutgers University
zuohui.fu@rutgers.edu

Yikun Xian
Dept of Computer Science
Rutgers University
yikun.xian@rutgers.edu

Yongfeng Zhang
Dept of Computer Science
Rutgers University
yongfeng.zhang@rutgers.edu

Yi Zhang
School of Engineering,
UC Santa Cruz
yiz@soe.ucsc.edu

ABSTRACT

Recent years have witnessed the emerging of conversational systems, including both physical devices and mobile-based applications. Both the research community and industry believe that conversational systems will have a major impact on human-computer interaction, and specifically, the IR/DM/RecSys communities have begun to explore Conversational Recommendation Systems. Conversational recommendation aims at finding or recommending the most relevant information (e.g., web pages, answers, movies, products) for users based on textual- or spoken-dialogs, through which users can communicate with the system more efficiently using natural language conversations. Due to users' constant need to look for information to support both work and daily life, conversational recommendation system will be one of the key techniques towards an intelligent web. The tutorial focuses on the foundations and algorithms for conversational recommendation, as well as their applications in real-world systems such as search engine, e-commerce and social networks. The tutorial aims at introducing and communicating conversational recommendation methods to the community, as well as gathering researchers and practitioners interested in this research direction for discussions, idea communications, and research promotions.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; • **Computing methodologies** → **Natural language processing**;

KEYWORDS

Conversational Recommendation; Dialog Systems

ACM Reference Format:

Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2021. WSDM 2021 Tutorial on Conversational Recommendation Systems. In *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM '21), March 8–12, 2021, Virtual Event, Israel*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3437963.3441661>

1 INTRODUCTION

Among many techniques that compose an intelligent Web, a Conversational System (such as Google Now, Apple Siri, Microsoft

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 the author(s) 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.

WSDM '21, March 8–12, 2021, Virtual Event, Israel

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8297-7/21/03...\$15.00
<https://doi.org/10.1145/3437963.3441661>

Cortana, and Amazon Alexa) is one that serves as the direct interactive portal for end-users, which is expected to revolutionize human-computer interaction. With recent progress on NLP and Internet of Things (IoT), such systems have also been deployed as physical devices, opening up more opportunities for applications in smart environments. Under this background, Conversational Recommendation System (CRS) has become an important direction to bridge conversational agents and personalized services [20, 29].

Though as an emerging research direction, the basic idea of conversational recommendation has been a core consideration for researchers in the field, from early-stage interactive collaborative filtering methods to assist cold-start users, to later reinforcement learning approaches for modeling user-dialog states, and to more recent deep learning and natural language processing approaches to conversational recommendation [1–3, 13, 19, 20, 23, 28, 29, 31–34].

This tutorial will introduce conversational recommendation from five main perspectives – dataset, method, evaluation, toolkit, and real-world systems. In particular:

Dataset: We introduce current datasets that facilitate conversational recommendation research including three types of existing datasets: fully-synthetic, semi-synthetic, and real-world datasets. Due to their difference in conversational pattern and data scale, different datasets facilitate different types of conversational recommendation research, and we introduce the differences and how to make best use of the datasets in detail. The problem definition of conversational recommendation also highly relies on the datasets and conversational pattern, as a result, we also formalize and introduce the problem definition of conversational recommendation in this part of the tutorial.

Method: In this part of the tutorial, we organize and introduce current methods and algorithms for conversational recommendation. We introduce both algorithmic approaches and user-analysis approaches to conversational recommendation. In particular, we carefully organize existing approaches into a two-dimensional taxonomy. One dimension is conversation paradigm, which can be either system-leading or user-leading. The second dimension is the method for conversation strategies. In this dimension, we first introduce early-stage interactive collaborative filtering methods, and then introduce more recent conversational recommendation approaches in terms of two broad types: dialog-state tracking and natural language understanding, which constitutes two important ingredients of practical conversational recommendation systems.

Evaluation: We introduce both offline and online evaluation methods for conversational recommendation. Evaluation methods can be different according to different conversation paradigms. In system-leading conversations, we introduce evaluation methods for question prediction, conversation efficiency, and recommendation accuracy. In user-leading conversations, we introduce evaluation methods for response accuracy and conversation efficiency. A scope

of important research exists in user-analysis and evaluation for conversational recommendation systems, and we also introduce these researches to facilitate better user-oriented evaluation of conversational recommendation systems.

Toolkit: We introduce existing open-source toolkits and platforms to facilitate conversational recommendation research, such as the Microsoft Conversational AI tools and bot builder SDK. This part of the tutorial includes a hands-on demo about quick building of a working conversational recommender system, so as to help attendees to gain hands-on experience in this research domain.

Real-world systems: Finally, we introduce conversational recommendation in current real-world conversational AI systems in industry, including but not limited to Google Assistant, Amazon Alexa, and Microsoft Cortana. In this part of the tutorial, we also introduce the programming interfaces supported in these systems to help attendees understand how to integrate their conversational recommendation methods into the systems.

Based on the above rich set of materials, this tutorial will greatly help researchers from both academia and industry who are interested in conversational recommendation.

2 THE TUTORIAL TEAM

Dr. Yongfeng Zhang and his group have been consistently working on conversational recommendation systems. Together with co-authors, their work on conversational recommendation [29] provided a standard System Ask–User Respond (SAUR) paradigm for conversational search and recommendation. The paradigm has helped researchers and practitioners to push the frontier in the field for different conversational search and recommendation scenarios. Their other research on conversational recommendation and information seeking includes [2, 14–18, 25–27]. Senior PhD students Zuohui Fu and Yikun Xian have been working on knowledge-graph reasoning approaches to conversational recommendation. A selected subset of their recent publications include [4–12, 21, 22, 24, 30]. Dr. Yi Zhang is a pioneer in conversational recommendation research. Her paper on conversational recommendation [20] is one of the first papers on dialog state controlling for conversational recommender systems, which inspired a sequence of research on this problem. Her experience as the CTO of a Conversational AI tech-company Rulai Inc.¹ further enhances the tutorial by providing hands-on industry experiences to our audiences.

3 BRIEF BIO OF ORGANIZERS

Zuohui Fu² is a Ph.D. student in the Department of Computer Science at Rutgers University, jointly advised by Prof. Yongfeng Zhang and Prof. Gerard de Melo. His research interests focus on solving the cutting-edge research problems over Information Retrieval, Natural Language Processing and Machine Learning, with the emphasis on semantic matching, dialogue systems, and decision making over knowledge graphs, etc. His recent research mainly aims at KG-assisted explainable, fair and conversational recommender system. His works have appeared in premier IR and artificial intelligence

conferences including AAAI, SIGIR, WWW, IJCAI, ACL, etc. Besides, he has served as the conference PC member in NeurIPS, AAAI, ACL, EMNLP, CIKM, etc., and the journal reviewer for ACM TOIS, etc. He has presented the Conversational Recommendation tutorial on RecSys 2020.

Yikun Xian³ is a Ph.D. student at the Department of Computer Science of Rutgers University, jointly advised by Prof. Yongfeng Zhang and Prof. Shan Muthukrishnan. His research interest broadly lies in the intersection of information retrieval and machine learning. His recent research include graph-based explainable decision-making problems with applications to explainable recommendation, knowledge-enhanced and explainable conversational recommendation, as well as interpretable comparative shopping. He mainly develops machine learning techniques such as graph-based reasoning, neural symbolic reasoning and deep reinforcement learning for information retrieval and recommender systems. He has served as PC member/reviewer in top computer science conference/journals such as NeurIPS, EMNLP, CIKM, SIGIR, ECML-PKDD, WWW, ECIR, RecSys, and ACM TOIS. He has presented the Conversational Recommendation tutorial on RecSys 2020.

Yongfeng Zhang⁴ is an Assistant Professor in the Department of Computer Science at Rutgers University (The State University of New Jersey). His research interest is in Information Retrieval, Recommender Systems, Machine Learning, and Internet Economics. In the previous he was a postdoc in the Center for Intelligent Information Retrieval (CIIR) at UMass Amherst, and did his PhD and BE in Computer Science at Tsinghua University, with a BS in Economics at Peking University. He is a Siebel Scholar of the class 2015, and a Baidu Scholar of the class 2014. He has been consistently working on conversational recommendation, explainable recommendation, and the intersection between the two subjects. His recent work on conversational recommendation models includes learning to ask for conversational recommendation, explainable conversational recommendation, and knowledge-enhanced conversational recommendation. He has presented the WWW2019/SIGIR2019/ICTIR2019 Tutorial on Explainable Recommendation and Search, and the RecSys 2020 Tutorial on Conversational Recommendation.

Yi Zhang⁵ is a professor in the School of Engineering, University of California Santa Cruz. Her research interests include large scale information retrieval, recommendation systems, internet advertising, data mining, natural language processing, and applied machine learning. She has published chapters, journal articles, and papers at top conferences in these areas, such as ACM SIGIR, WWW, CIKM, IEEE ICDM, ICML, COLINGS, HLT. She received NSF Faculty Early Career Award in 2010, an Air Force Research Young Investigator Award in 2008, the Best Paper Award at ACM SIGIR in 2002, and several other awards. Her Information Retrieval and Knowledge Management Lab is doing research sponsored by several government agencies and companies (Microsoft, Yahoo, Google, NEC, Bosch, Nokia etc.). She has served as a consultant or technical advisor for companies. She regularly serves on the program committees

¹<https://rul.ai/>

²<https://zuohuif.github.io/>

³<https://orcax.github.io/>

⁴<http://yongfeng.me>

⁵<https://users.soe.ucsc.edu/~yiz/>

of the very best conferences in her research areas. She has served as area chair or senior PC member at ACM SIGIR, EMNLP, and ACM Recommender Systems. She has served as conference co-chair in charge of Information Retrieval area at the ACM Conference on Information and Knowledge Management, and tutorial chair for ACM SIGIR. She is serving as an associate editor for ACM Transaction on Information Systems. Dr. Zhang received her Ph.D. from School of Computer Science at CMU, specializing in Language and Information Technologies. She has presented the Conversational Recommendation tutorial on RecSys 2020.

4 FORMAT AND INTERACTION STYLES

The tutorial is delivered as a full tutorial (3 hours plus break). One and a half hours for the dataset and method section with QA, and another 1.5 hour for the evaluation, toolkit and real-world system section with QA. At the end of the tutorial, we aggregate and wrap up different conversational methods, together with a final QA session, so as to encourage idea dissemination and research communication among the audiences.

5 AUDIENCE

The tutorial mainly targets on information retrieval, data mining and especially recommendation system researchers and practitioners. Since we also introduce how recent NLP and Knowledge base techniques will help conversational recommendation, it is also relevant to NLP, Semantic, and Knowledge researchers. Since we also introduce how conversational recommendation are applied in commercial real-world systems such as e-commerce, search engine, and social networks, it is also helpful to industry researchers and practitioners from different areas. For prerequisite, basic understandings of information retrieval and recommendation system knowledge will be preferred, but the tutorial will introduce the basic concepts for better audience engagement.

6 TUTORIAL MATERIALS

Tutorial materials such as the slides and video recordings are provided on the tutorial website (<http://conversational-recsys.github.io>) to the audiences and the general public.

REFERENCES

- [1] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W Bruce Croft. 2019. Asking clarifying questions in open-domain information-seeking conversations. In *SIGIR*. 475–484.
- [2] Keping Bi, Qingyao Ai, Yongfeng Zhang, and W Bruce Croft. 2019. Conversational product search based on negative feedback. In *CIKM*. 359–368.
- [3] Zhongxia Chen, Xiting Wang, Xing Xie, Mehul Parsana, Akshay Soni, Xiang Ao, and Enhong Chen. 2020. Towards Explainable Conversational Recommendation. *IJCAI* (2020).
- [4] Zuohui Fu, Yikun Xian, Ruoyuan Gao, Jieyu Zhao, Qiaoying Huang, Yingqiang Ge, Shuyuan Xu, Shijie Geng, Chirag Shah, Yongfeng Zhang, and Gerard de Melo. 2020. Fairness-Aware Explainable Recommendation over Knowledge Graphs. *SIGIR* (2020).
- [5] Zuohui Fu, Yikun Xian, Shijie Geng, Yingqiang Ge, Yuting Wang, Xin Dong, Guang Wang, and Gerard de Melo. 2020. ABSent: Cross-Lingual Sentence Representation Mapping with Bidirectional GANs. In *AAAI*.
- [6] Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2020. Tutorial on Conversational Recommendation Systems. In *RecSys*. 751–753.
- [7] Zuohui Fu, Yikun Xian, Yaxin Zhu, Yongfeng Zhang, and Gerard de Melo. 2020. COOKIE: A Dataset for Conversational Recommendation over Knowledge Graphs in E-commerce. *arXiv preprint arXiv:2008.09237* (2020).
- [8] Yingqiang Ge, Shuchang Liu, Ruoyuan Gao, Yikun Xian, Yunqi Li, Xiangyu Zhao, Changhua Pei, Fei Sun, Junfeng Ge, Wenwu Ou, and Yongfeng Zhang. 2021. Towards Long-term Fairness in Recommendation. In *WSDM*.
- [9] Yingqiang Ge, Shuyuan Xu, Shuchang Liu, Zuohui Fu, Fei Sun, and Yongfeng Zhang. 2020. Learning Personalized Risk Preferences for Recommendation. In *SIGIR*. 409–418.
- [10] Yingqiang Ge, Shuyuan Xu, Shuchang Liu, Shijie Geng, Zuohui Fu, and Yongfeng Zhang. 2019. Maximizing marginal utility per dollar for economic recommendation. In *The World Wide Web Conference*. 2757–2763.
- [11] Yunqi Li, Shuyuan Xu, Bo Liu, Zuohui Fu, Shuchang Liu, Xu Chen, and Yongfeng Zhang. 2020. Discrete Knowledge Graph Embedding based on Discrete Optimization. In *Proceedings of the AAAI-20 Workshop on Knowledge Discovery from Unstructured Data in Financial Services*.
- [12] Zhenyu Liao, Yikun Xian, Xiao Yang, Qinpei Zhao, Chenxi Zhang, and Jiangfeng Li. 2018. TSCSet: A crowdsourced time-sync comment dataset for exploration of user experience improvement. In *IUI*. 641–652.
- [13] Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. *arXiv preprint arXiv:2005.03954* (2020).
- [14] Chen Qu, Liu Yang, W Bruce Croft, Falk Scholer, and Yongfeng Zhang. 2019. Answer interaction in non-factoid question answering systems. In *CHIIR*.
- [15] Chen Qu, Liu Yang, W Bruce Croft, Johanne R Trippas, Yongfeng Zhang, and Minghui Qiu. 2018. Analyzing and characterizing user intent in information-seeking conversations. In *SIGIR*. 989–992.
- [16] Chen Qu, Liu Yang, W Bruce Croft, Yongfeng Zhang, Johanne R Trippas, and Minghui Qiu. 2019. User intent prediction in information-seeking conversations. In *ACM CHIIR*.
- [17] Chen Qu, Liu Yang, Minghui Qiu, W Bruce Croft, Yongfeng Zhang, and Mohit Iyyer. 2019. BERT with History Answer Embedding for Conversational Question Answering. In *SIGIR*. 1133–1136.
- [18] Chen Qu, Liu Yang, Minghui Qiu, Yongfeng Zhang, Cen Chen, W Bruce Croft, and Mohit Iyyer. 2019. Attentive History Selection for Conversational Question Answering. In *CIKM*. 1391–1400.
- [19] Corbin Rosset, Chenyan Xiong, Xia Song, Daniel Campos, Nick Craswell, Saurabh Tiwary, and Paul Bennett. 2020. Leading Conversational Search by Suggesting Useful Questions. In *Proceedings of The Web Conference 2020*. 1160–1170.
- [20] Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *SIGIR*. 235–244.
- [21] Yikun Xian, Zuohui Fu, S Muthukrishnan, Gerard De Melo, and Yongfeng Zhang. 2019. Reinforcement knowledge graph reasoning for explainable recommendation. In *SIGIR*. 285–294.
- [22] Yikun Xian, Zuohui Fu, Handong Zhao, Yingqiang Ge, Xu Chen, Qiaoying Huang, Shijie Geng, Zhou Qin, Gerard de Melo, S. Muthukrishnan, and Yongfeng Zhang. 2020. CAFE: Coarse-to-Fine Neural Symbolic Reasoning for Explainable Recommendation. *CIKM* (2020).
- [23] Hu Xu, Seungwhan Moon, Honglei Liu, Bing Liu, Pararth Shah, and Philip S Yu. 2020. User Memory Reasoning for Conversational Recommendation. *arXiv preprint arXiv:2006.00184* (2020).
- [24] Shuyuan Xu, Yunqi Li, Shuchang Liu, Zuohui Fu, and Yongfeng Zhang. 2020. Learning Post-Hoc Causal Explanations for Recommendation. *arXiv preprint arXiv:2006.16977* (2020).
- [25] Liu Yang, Minghui Qiu, Chen Qu, Cen Chen, Jiafeng Guo, Yongfeng Zhang, W Bruce Croft, and Haiqing Chen. 2020. IART: Intent-aware Response Ranking with Transformers in Information-seeking Conversation Systems. In *WWW*.
- [26] Liu Yang, Minghui Qiu, Chen Qu, Jiafeng Guo, Yongfeng Zhang, W Bruce Croft, Jun Huang, and Haiqing Chen. 2018. Response ranking with deep matching networks and external knowledge in information-seeking conversation systems. In *SIGIR*. 245–254.
- [27] Liu Yang, Hamed Zamani, Yongfeng Zhang, Jiafeng Guo, and W Bruce Croft. 2017. Neural matching models for question retrieval and next question prediction in conversation. *arXiv preprint arXiv:1707.05409* (2017).
- [28] Hamed Zamani, Susan Dumais, Nick Craswell, Paul Bennett, and Gord Lueck. 2020. Generating clarifying questions for information retrieval. In *Proceedings of The Web Conference 2020*. 418–428.
- [29] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *CIKM*. 177–186.
- [30] Honglu Zhou, Shuyuan Xu, Zuohui Fu, Gerard de Melo, Yongfeng Zhang, and Mubbasir Kapadia. 2020. HID: Hierarchical Multiscale Representation Learning for Information Diffusion. *IJCAI* (2020).
- [31] Kun Zhou, Wayne Xin Zhao, Shuqing Bian, Yuanhang Zhou, Ji-Rong Wen, and Jingsong Yu. 2020. Improving Conversational Recommender Systems via Knowledge Graph based Semantic Fusion. *KDD* (2020).
- [32] Jie Zou, Yifan Chen, and Evangelos Kanoulas. 2020. Towards Question-based Recommender Systems. *SIGIR* (2020).
- [33] Jie Zou and Evangelos Kanoulas. 2019. Learning to Ask: Question-based Sequential Bayesian Product Search. In *CIKM*. 369–378.
- [34] Jie Zou and Evangelos Kanoulas. 2020. Towards Question-Based High-Recall Information Retrieval: Locating the Last Few Relevant Documents for Technology Assisted Reviews. *ACM Transactions on Information Systems (TOIS)* (2020).