

IUI 2021 Tutorial on Conversational Recommendation Systems

Zuohui Fu¹, Yikun Xian¹, Yongfeng Zhang¹, Yi Zhang²

¹Rutgers University, NJ ²UC Santa Cruz, CA

United States

{zuohui.fu,yikun.xian,yongfeng.zhang}@rutgers.edu,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 CHI/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.

KEYWORDS

Conversational Recommendation; Dialog Systems; Natural Language Processing; Intelligent Interface; Human-in-the-Loop AI

ACM Reference Format:

Zuohui Fu¹, Yikun Xian¹, Yongfeng Zhang¹, Yi Zhang². 2021. IUI 2021 Tutorial on Conversational Recommendation Systems. In *26th International Conference on Intelligent User Interfaces (IUI '21 Companion)*, April 14–17, 2021, College Station, TX, USA. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3397482.3450621>

1 MOTIVATION, TOPICS AND RELEVANCE

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

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.

IUI '21 Companion, April 14–17, 2021, College Station, TX, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8018-8/21/04...\$15.00

<https://doi.org/10.1145/3397482.3450621>

physical devices, opening up more opportunities for applications in smart environments.

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.

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 [35] 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 such as [2, 6, 16, 22, 29, 37, 38]. Their other research on conversational recommendation and information seeking includes [2, 9, 17–21, 31–33]. 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 [7, 8, 14, 27, 28]. The group also has research experience on explainable recommendation [1, 5, 7, 26–28, 34, 36], natural language generation [3, 13], neural logic/symbolic reasoning [4, 23, 24] and economic modeling of user behaviors [10–12, 15, 30], which help to strengthen the tutorial from broader perspectives. Dr. Yi Zhang is a pioneer in conversational recommendation research. Her paper on conversational recommendation [25] is one of the first papers on dialog state controlling for conversational recommender systems. Her experience as the CTO of a Conversational AI tech-company Rulai Inc. further enhances the tutorial by providing hands-on industry experiences to our audiences.

3 FORMAT AND INTERACTION STYLES

The tutorial is delivered as a half-day tutorial. We will introduce the background and problem formalization (by Yongfeng Zhang), dataset and evaluation (by Yikun Xian), conversational recommendation models (by Zuohui Fu), as well as toolkit and real-world systems (by Yi Zhang). At the end of the tutorial, we will 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. The tutorial targets on CHI, IR, NLP, data mining and especially recommender system researchers and practitioners.

4 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] Q. Ai, V. Azizi, X. Chen, and Y. Zhang. 2018. Learning heterogeneous knowledge base embeddings for explainable recommendation. *Algorithms* 11, 9 (2018), 137.
- [2] Keping Bi, Qingyao Ai, Yongfeng Zhang, and W Bruce Croft. 2019. Conversational product search based on negative feedback. In *CIKM*. 359–368.
- [3] H. Chen, X. Chen, S. Shi, and Y. Zhang. 2019. Generate natural language explanations for recommendation. In *Proceedings of the 2nd EARS Workshop*.
- [4] Hanxiong Chen, Shaoyun Shi, Yunqi Li, and Yongfeng Zhang. 2020. Neural Collaborative Reasoning. *arXiv preprint arXiv:2005.08129* (2020).
- [5] Xu Chen, Yongfeng Zhang, and Zheng Qin. 2019. Dynamic explainable recommendation based on neural attentive models. In *AAAI*, Vol. 33. 53–60.
- [6] Zhongxia Chen, Xiting Wang, Xing Xie, Mehul Parsana, Akshay Soni, et al. 2020. Towards Explainable Conversational Recommendation. *IJCAI* (2020).
- [7] Z. Fu, Y. Xian, R. Gao, J. Zhao, Q. Huang, Y. Ge, S. Xu, et al. 2020. Fairness-Aware Explainable Recommendation over Knowledge Graphs. *SIGIR* (2020).
- [8] Zuohui Fu, Yikun Xian, Shijie Geng, Yingqiang Ge, et al. 2020. ABSent: Cross-Lingual Sentence Representation Mapping with Bidirectional GANs. In *AAAI*.
- [9] Z. Fu, Y. Xian, Y. Zhu, Y. Zhang, et al. 2020. COOKIE: A Dataset for Conversational Recommendation over Knowledge Graphs in E-commerce. *arXiv* (2020).
- [10] Y. Ge, S. Xu, S. Liu, Z. Fu, F. Sun, and Y. Zhang. 2020. Learning Personalized Risk Preferences for Recommendation. In *SIGIR*.
- [11] Y. Ge, S. Xu, S. Liu, S. Geng, Z. Fu, and Y. Zhang. 2019. Maximizing marginal utility per dollar for economic recommendation. In *WWW*.
- [12] Y. Ge, S. Zhao, H. Zhou, C. Pei, F. Sun, W. Ou, and Y. Zhang. 2020. Understanding echo chambers in e-commerce recommender systems. In *SIGIR*.
- [13] Lei Li, Yongfeng Zhang, and Li Chen. 2020. Generate neural template explanations for recommendation. In *CIKM*. 755–764.
- [14] Z. Liao, Y. Xian, X. Yang, Q. Zhao, et al. 2018. TSCSet: A crowdsourced time-sync comment dataset for exploration of user experience improvement. In *IUI*.
- [15] X. Lin, H. Chen, C. Pei, F. Sun, X. Xiao, et al. 2019. A pareto-efficient algorithm for multiple objective optimization in e-commerce recommendation. In *RecSys*.
- [16] Z. Liu, H. Wang, Z. Niu, H. Wu, W. Che, and T. Liu. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. *arXiv* (2020).
- [17] Chen Qu, Liu Yang, W Bruce Croft, Falk Scholer, and Yongfeng Zhang. 2019. Answer interaction in non-factoid question answering systems. In *CHIIR*.
- [18] C. Qu, L. Yang, W. B. Croft, J. R. Trippas, Y. Zhang, and M. Qiu. 2018. Analyzing and characterizing user intent in information-seeking conversations. In *SIGIR*.
- [19] C. Qu, L. Yang, W. B. Croft, Y. Zhang, J. R. Trippas, and M. Qiu. 2019. User intent prediction in information-seeking conversations. In *CHIIR*.
- [20] C. Qu, L. Yang, M. Qiu, W. B. Croft, Y. Zhang, and M. Iyyer. 2019. BERT with History Answer Embedding for Conversational Question Answering. In *SIGIR*.
- [21] C. Qu, L. Yang, M. Qiu, Y. Zhang, C. Chen, W. B. Croft, et al. 2019. Attentive History Selection for Conversational Question Answering. In *CIKM*. 1391–1400.
- [22] C. Rosset, C. Xiong, X. Song, D. Campos, N. Craswell, S. Tiwary, and P. Bennett. 2020. Leading Conversational Search by Suggesting Useful Questions. In *WWW*.
- [23] Shaoyun Shi, Hanxiong Chen, Weizhi Ma, Jiaxin Mao, Min Zhang, and Yongfeng Zhang. 2020. Neural Logic Reasoning. In *CIKM*. 1365–1374.
- [24] Shaoyun Shi, Hanxiong Chen, Min Zhang, and Yongfeng Zhang. 2019. Neural logic networks. *arXiv preprint arXiv:1910.08629* (2019).
- [25] Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *SIGIR*.
- [26] Pengfei Wang, Hanxiong Chen, Yadong Zhu, Huawei Shen, and Yongfeng Zhang. 2019. Unified Collaborative Filtering over Graph Embeddings. In *SIGIR*. 155–164.
- [27] Y. Xian, Z. Fu, S. Muthukrishnan, G. De Melo, and Y. Zhang. 2019. Reinforcement knowledge graph reasoning for explainable recommendation. In *SIGIR*.
- [28] Y. Xian, Z. Fu, H. Zhao, Y. Ge, X. Chen, et al. 2020. CAFE: Coarse-to-Fine Neural Symbolic Reasoning for Explainable Recommendation. *CIKM* (2020).
- [29] H. Xu, S. Moon, H. Liu, B. Liu, P. Shah, and P. S. Yu. 2020. User Memory Reasoning for Conversational Recommendation. *arXiv* (2020).
- [30] Zhichao Xu, Yi Han, Yongfeng Zhang, and Qingyao Ai. 2020. E-commerce Recommendation with Weighted Expected Utility. In *CIKM*.
- [31] L. Yang, M. Qiu, C. Qu, et al. 2018. Response ranking with deep matching networks and external knowledge in information-seeking conversation systems. In *SIGIR*.
- [32] L. Yang, M. Qiu, C. Qu, C. Chen, et al. 2020. IART: Intent-aware Response Ranking with Transformers in Information-seeking Conversation Systems. In *WWW*.
- [33] L. Yang, H. Zamani, Y. Zhang, J. Guo, et al. 2017. Neural matching models for question retrieval and next question prediction in conversation. *arXiv* (2017).
- [34] Yongfeng Zhang and Xu Chen. 2020. Explainable Recommendation: A Survey and New Perspectives. *Foundations and Trends in Information Retrieval* (2020).
- [35] Y. Zhang, X. Chen, Q. Ai, L. Yang, and W. B. Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *CIKM*.
- [36] Y. Zhang, G. Lai, M. Zhang, Y. Zhang, Y. Liu, et al. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In *SIGIR*.
- [37] K. Zhou, W. X. Zhao, S. Bian, et al. 2020. Improving Conversational Recommender Systems via Knowledge Graph based Semantic Fusion. *KDD* (2020).
- [38] Jie Zou, Yifan Chen, and Evangelos Kanoulas. 2020. Towards Question-based Recommender Systems. *SIGIR* (2020).