Proactive Conversational Agents

Lizi Liao
Singapore Management University
lzliao@smu.edu.sg

Grace Hui Yang
Georgetown University
grace.yang@georgetown.edu

Chirag Shah
University of Washington
chirags@uw.edu

ABSTRACT

Conversational agents, or commonly known as dialogue systems, have gained escalating popularity in recent years. Their widespread applications support conversational interactions with users and accomplishing various tasks as personal assistants. However, one key weakness in existing conversational agents is that they only learn to passively answer user queries via training on pre-collected and manually-labeled data. Such passiveness makes the interaction modeling and system-building process relatively easier, but it largely hinders the possibility of being human-like hence lowering the user engagement level. In this tutorial, we introduce and discuss methods to equip conversational agents with the ability to interact with end users in a more proactive way. This three-hour tutorial is divided into three parts and includes two interactive exercises. It reviews and presents recent advancements on the topic, focusing on automatically expanding ontology space, actively driving conversation by asking questions or strategically shifting topics, and retrospectively conducting response quality control.

KEYWORDS

Proactive conversational agents, task-oriented dialogue systems, conversational search, conversational AI

ACM Reference Format:
Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. Proactive Conversational Agents. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining (WSDM '23), February 27-March 3, 2023, Singapore, Singapore.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3539597.3572724

1 MOTIVATION & BRIEF OUTLINE

Building conversational agents capable of understanding human language and making proper responses to humans is a long-cherished goal of Artificial Intelligence (AI). Such agents can be broadly categorized into two main types: (1) Chit-chat systems that are designed to engage users and provide emotional support. It conducts chit-chat type of conversations on open topics without having any specific goal to complete, and (2) Task-oriented dialogue systems that are designed to assist users in completing certain tasks. Currently, most of the popular personal assistants, such as Amazon Alexa, Apple Siri, Google Home, and Microsoft Cortana, are task-oriented. They are primarily designed to take natural language verbal instructions or commands, understand them and translate them into actions to be executed by the underlying application systems.

Despite the increasing attention and widespread adaptations, most existing conversational agents can only perform pre-defined actions or offer factual information in response to a user's command or query. They heavily rely on existing training conversations and exploit knowledge associated with these for exchanging information, recommending something, and completing tasks [21, 25, 63]. Such systems usually passively generate responses to questions instead of leading the conversation or asking questions [5, 50, 51]. This reactive, passive style of response constrains the types of conversations one could have. For example, situations involving exploratory search and complex decision-making where a conversation with both sides actively engaging is desired.

In the past a few years, researchers in multiple fields, including natural language processing (NLP), dialogue systems, conversational search, conversational recommender systems, and machine learning [8, 12, 14, 35, 62, 64] have started to address the issue to enable conversational agents the ability for two-way, proactive conversations [30, 32, 52, 60, 65]. Various aspects or angles have been touched, such as proposing new intents and conversational slots via ontology expansion and actively analyzing failures in unseen situations [17, 18, 23, 27–29, 34, 40, 41, 48, 49, 59, 61], learning to raise questions to move the conversation forward [22, 36, 66, 67], leveraging task information and domain knowledge to guide purposeful topic shifting and exploration [26, 38, 39, 42, 56], and controlling quality of the response generation [2, 11, 33, 43, 46] as well as improving evaluation [16, 44] etc. In this tutorial, we review the literature on this emerging research topic, and introduce and discuss methods in a structured way to equip conversational agents with the ability to proactively interact with end users in performing various tasks.

This three-hour tutorial will be divided into three parts and two interactive exercises. The first part of the tutorial focuses on introducing the paradigm of proactive conversation agents and discusses system architecture design, the differences and challenges in realizing the proactive conversational AI paradigm. It includes an interactive exercise to engage the audience while emphasizing the important issues to address. In the second part, we present recent advancements on the topic, focusing on automatically expanding ontology space and how to drive conversation by asking questions actively. In the third part, we review and discuss methods that support strategic planning of topic shifting and retrospectively conducting response quality control. At the end of the lectures, we arrange a session of interactive exercises for the attendees to chat with a few conversational agents supported by major variations of the methods we show in the tutorial and test their own fixes to solve some of the issues that will be present. Finally, we conclude

WSDM '23, February 27-March 3, 2023, Singapore, Singapore

ACM ISBN 978-1-4503-9407-9/23/02…$15.00
https://doi.org/10.1145/3539597.3572724

with a discussion on the scopes for proactive conversational agents and present some open challenges for future research.

2 INTENDED AUDIENCE & PREREQUISITE

The intended audience includes researchers, graduate or senior undergraduate students, and industrial practitioners interested in conversational search, conversational recommendation, and task-oriented dialogue systems. This timely tutorial will particularly benefit those who are building or intend to build conversational agents for being proactive is becoming necessary for next-generation conversational agents. The tutorial will assume the audience's basic knowledge and experience in machine learning, deep learning, information retrieval, and NLP.

3 RELEVANCE TO THE COMMUNITY

We observe a steadily increasing interest in conversational agents in recent IR, AI, and NLP conferences. For instance, there is a Conversational Assistant Track in TREC 2019 to 2022[8], and the primer IR conference, SIGIR, has recognized 'Dialogue and Interactive systems' to be one of its major research topics. Our tutorial focuses on building next-generation conversational agents that can interact with end users proactively and contribute more to the shared ongoing conversation. With the wide use of customer service conversational agents, task completion conversational agents, and virtual personal assistants *etc.* in our daily life, there has been urgent demand to make these conversational agents more proactive. As more and more usage of conversational interfaces in downstream tasks becomes prevalent in both commercial and daily use cases, we believe the content of this tutorial would be interesting, beneficial and important to both IR researchers and practitioners.

Similar Tutorials in Related Conferences: Several tutorials on general dialogue systems and conversational recommendation have been given in recent related conferences, such as *Conversational Information Seeking* [7] in SIGIR'22 and *Conversational Recommendation Systems* [13] in RecSys'20 and in WSDM'21. However, none of them had focused on the proactive aspect of conversational agents. With a completely different and new perspective, our tutorial expands the horizon of conversational agents in a timely fashion and presents the latest advances on topics that are complementary yet relevant to existing mainstream approaches (e.g., topic shifting, ontology expansion, out-of-distribution machine learning, inappropriate dialogue management, and conservative reinforcement learning). We are unaware of any other proactive conversational agents tutorials in WSDM or related conferences.

4 FORMAT AND DETAILED SCHEDULE

Our tutorial will mainly be lectures, with interactive exercises created for both in-person and online attendees. We will also maintain an online forum to host question-answering sessions and encourage discussions during and after the tutorial.

- Schedule of the Tutorial [3hrs]

Part One:

(1) **Introduction———————————————————[15 mins]**
  (a) Dialogue & Interactive Systems: Background
  (b) State-of-the-art conversational agents: Weaknesses
  (c) Motivation & Aim of this tutorial

(2) **Interactive Exercise-1—————————————— [15 mins]**

(3) **Proactive Conv. Agents: Overview————————-[25 mins]**
  (a) Existing conversational agents architectures
  (b) Proactive extensions
  (c) Applications: Search, recommend, and more [15, 54]
  (d) Evaluation: Metrics, procedures, and challenges [16, 44]

Part Two:

(4) **Proactive Ontology Expansion ——————————[25 mins]**
  (a) New intent discovery [28, 61]
  (b) Novel slot induction [24, 49, 53, 55, 59]
  (c) Reflecting experiences for unseen situations [58]

(5) **Learning to Ask————————————————[25 mins]**
  (a) Mixed initiatives [19, 47]
  (b) Learning to ask [22, 36, 66, 67]
  (c) Question selection and generation [3, 9, 57]

Part Three:

(6) **Purposeful Topic Shifting——————————-[25 mins]**
  (a) Target-guided topic shifting [26, 42, 56]
  (b) Counterfactual data augmentation [4, 19, 43, 45]
  (c) Out-of-distribution label generation [1, 31]

(7) **Response Quality Control———————————[25 mins]**
  (a) Types of inappropriate responses [10]
  (b) Language detoxification [2, 11, 33]
  (c) Conservative reinforcement learning [4, 6, 20, 37, 43, 46]

(8) **Interactive Exercise-2 ————————————— [15 mins]**

Ending:

(9) **Concluding Remarks ——————————————-[10 mins]**
  (a) Summary
  (b) Challenges & Future directions

As listed in the scheduled above, we will have two interactive exercises. In the first one that takes place during Part One, we will instruct the participants to pair up. One of them in each pair will act as a user and the other one will act as an agent. We will then assign the users a task that requires gathering information and making a decision. An example task would be planning a trip to Singapore. The person acting as the agent can respond to a question coming from the user by looking up relevant information (e.g., from a search engine or a website) and constructing a short answer, similar to a typical conversational assistant will do. We will spend about 5 minutes setting up the exercise, 5 minutes having the participants do the role-play, and 5 minutes to come back and share what we discovered. This exercise will do two things: (1) it will allow us to see some of the limitations of the current systems that will set the stage for the next part of the tutorial; and (2) it will set up a baseline for the next interactive exercise.

The second interactive exercise will take place in Part Three, almost toward the end. By this time, the participants will have learned various methods that can allow a conversational agent to be proactive. We will now repeat the exercise from before, but this time asking the persons playing the agent to use appropriate technique (e.g., ontology expansion, topic shifting) while constructing their responses. Once again, we will come back and share our experiences. This will allow everyone to not only practice what they learned in the tutorial, albeit in a pseudo-system fashion, but also start thinking beyond the specific techniques to also account for human experiences in a proactive conversation.

- Reading List

Please see the References section for the Reading list.

5 SUPPORT MATERIALS

We will release our presentation slides, exercise questions and answers, and reference lists to the attendees. We also agree to allow video recording of the tutorial and make the videos available to the tutorial attendees. During and after the tutorial, we will collect questions from the attendees, post our answers to their questions, and maintain a discussion forum during the entire conference.

6 PARTICIPATION AND ENGAGEMENT

Weeks before the conference, we will promote the tutorial on social media and on professional email lists (e.g., the SIGIR mailing list). We will set up and maintain a tutorial website and promote the content there, too. We plan to contact colleagues in multiple research areas, including NLP, AI, Speech, Dialogue Systems, ML, Recommender systems, and IR, to forward the news about the tutorial. During the tutorial, we will have multiple sessions of interactive exercises with the attendees. We will also actively solicit questions from the attendees and post our answers to the questions in a forum (e.g. piazza.com) that will be created specifically for this tutorial. We will maintain the forum during the entire conference. The goal is to create the discussion and keep the communications going. For vivid presentations and hands-on experiences, conversational devices, such as Amazon Alexa and Apple Siri, will be installed in the tutorial room. They will be used in some examples in our slides and in interactive exercises with the attendees.

7 CONTACT INFORMATION OF PRESENTERS

(1) Lizi Liao (Main Contact), Singapore Management University
lzliao@smu.edu.sg
(2) Grace Hui Yang, Georgetown University
grace.yang@georgetown.edu
(3) Chirag Shah, University of Washington
chirags@uw.edu

8 BIOS OF PRESENTERS

Dr. Lizi Liao is Assistant Professor of Computer Science at Singapore Management University, leading the CoAgent group. She obtained her Ph.D. from National University of Singapore in 2019. Dr. Liao's research interests center on task-oriented dialogues, proactive conversational agents, and multimodal conversational search and recommendation as the application target. She publishes regularly in top AI and Data Science conferences and journals such as SIGIR, WWW, ACM MM, TKDE, TACL etc. One of her work was nominated in the ACM MM Best Paper Final List in 2018. She serves as (senior) PC member of these prestigious conferences and reviewer of the top journals. She also serves as organizing committee members of ACM MM 2019, NLPCC 2022 and WSDM 2023, while chairs sessions at KDD 2021 and SIGIR 2022.

Dr. Grace Hui Yang is Associate Professor of Computer Science at Georgetown University, Washington D.C., U.S., leading the InfoSense (Information Retrieval and Sense-Making) group. Dr. Yang obtained her Ph.D. from Carnegie Mellon University in 2011. Her current research interests include conversational agents, deep reinforcement learning, and information retrieval. Dr. Yang is a recipient of the U.S. National Science Foundation Faculty Early Career Development Program (CAREER) Award. She has served on the organizing or program committees in recent conferences, such as SIGIR, ECIR, ACL, AAAI, ICTIR, CIKM, WSDM, and WWW, and on the editorial boards of ACM TOIS (2022-) and Information Retrieval Journal (2014-2017). Dr. Yang has experience conducting tutorials in the WSDM conferences – she delivered tutorials on *Dynamic Information Retrieval* at SIGIR'14 and WSDM'15, and on *Privacy-Preserving IR* at ICTIR'17 and WSDM'18. Dr. Yang also taught in the first ACM SIGIR AFIRM Summer School in 2019.

Dr. Chirag Shah is Professor at University of Washington in Seattle. He has taught undergraduate and graduate courses in IR, HCI, Data Science, and Machine Learning. He has also taught several courses and tutorials on topics related to web search and recommender systems at different international places, including at conferences such as RecSys, SIGIR, and WSDM, Russian Summer School on Information Retrieval (RuSSIR), and Asian Summer School in Information Access (ASSIA). He has developed Coursera course on Social Media Data Analytics, and a SAGE course on Machine Learning for Data Science. He has authored multiple books, including a textbook on Data Science (2020) and soon to be published textbook on Machine Learning.

REFERENCES

[1] Christina Baek, Yiding Jiang, Aditi Raghunathan, and Zico Kolter. 2022. Agreement-on-the-Line: Predicting the Performance of Neural Networks under Distribution Shift.
[2] Ashutosh Baheti, Maarten Sap, Alan Ritter, and Mark Riedl. 2021. Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts. In *EMNLP*. 4846–4862.
[3] Keping Bi, Qingyao Ai, and W. Bruce Croft. 2021. Asking Clarifying Questions Based on Negative Feedback in Conversational Search. In *SIGIR'21*. ACM.
[4] Limin Chen, Zhiwen Tang, and Grace Hui Yang. 2020. Balancing Reinforcement Learning Training Experiences in Interactive Information Retrieval. In *SIGIR*. 1525–1528.
[5] Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. 2022. SalesBot: Transitioning from Chit-Chat to Task-Oriented Dialogues. In *ACL*. 6143–6158.
[6] Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. 2018. Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models. In *NeurIPS*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31.
[7] Jeffrey Dalton, Sophie Fischer, Paul Owoicho, Filip Radlinski, Federico Rossetto, Johanne R. Trippas, and Hamed Zamani. 2022. Conversational Information Seeking: Theory and Application. In *SIGIR'22*.
[8] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2021. TREC CAsT 2021: The Conversational Assistance Track Overview. In *TREC*.

[9] Bidyut Das, Mukta Majumder, Arif Ahmed Sekh, and Santanu Phadikar. 2022. Automatic question generation and answer assessment for subjective examination. *Cognitive Systems Research* 72 (2022), 14–22.

[10] Emily Dinan, Gavin Abercrombie, A. Bergman, Shannon Spruit, Dirk Hovy, Y-Lan Boureau, and Verena Rieser. 2022. SafetyKit: First Aid for Measuring Safety in Open-domain Conversational Systems. In *ACL*. 4113–4133.

[11] Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. Build it Break it Fix it for Dialogue Safety: Robustness from Adversarial Human Attack. In *EMNLP*.

[12] Zuohui Fu, Yikun Xian, Shijie Geng, Gerard De Melo, and Yongfeng Zhang. 2021. Popcorn: Human-in-the-loop Popularity Debiasing in Conversational Recommender Systems. In *CIKM*. 494–503.

[13] Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2020. Tutorial on Conversational Recommendation Systems. In *RecSys*.

[14] Prakhar Gupta, Jeffrey P Bigham, Yulia Tsvetkov, and Amy Pavel. 2021. Controlling Dialogue Generation with Semantic Exemplars. In *NAACL*. 3018–3029.

[15] Yingxu He, Lizi Liao, Zheng Zhang, and Tat-Seng Chua. 2021. Towards Enriching Responses with Crowd-sourced Knowledge for Task-oriented Dialogue. In *ACM MM Workshop on MuCAI*. 3–11.

[16] Chathra Hendahewa and Chirag Shah. 2017. Evaluating user search trails in exploratory search tasks. *IP&M* (2017), 905–922.

[17] Chengyu Huang, Zheng Zhang, Hao Fei, and Lizi Liao. 2022. Conversation Disentanglement with Bi-Level Contrastive Learning. In *EMNLP Findings*. To appear.

[18] Vojtěch Hudeček, Ondřej Dušek, and Zhou Yu. 2021. Discovering dialogue slots with weak supervision. In *ACL-IJCNLP*. 2430–2442.

[19] Rachna Konigari, Saurabh Ramola, Vijay Vardhan Alluri, and Manish Shrivastava. 2021. Topic Shift Detection for Mixed Initiative Response. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*.

[20] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative Q-Learning for Offline Reinforcement Learning. *CoRR* abs/2006.04779 (2020).

[21] Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. 2018. Towards Deep Conversational Recommendations. In *NeurIPS*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.).

[22] Zixuan Li, Lizi Liao, and Tat-Seng Chua. 2022. Learning to Ask Critical Questions for Assisting Product Search. In *SIGIR Workshop on eComm*.

[23] Lizi Liao, Le Hong Long, Yunshan Ma, Wenqiang Lei, and Tat-Seng Chua. 2021. Dialogue state tracking with incremental reasoning. *TACL* 9 (2021), 557–569.

[24] Lizi Liao, Le Hong Long, Zheng Zhang, Minlie Huang, and Tat-Seng Chua. 2021. MMConv: an environment for multimodal conversational search across multiple domains. In *SIGIR*. 675–684.

[25] Lizi Liao, Yunshan Ma, Xiangnan He, Richang Hong, and Tat-seng Chua. 2018. Knowledge-aware multimodal dialogue systems. In *ACM MM*. 801–809.

[26] L. Liao, R. Takanobu, Y. Ma, X. Yang, M. Huang, and T. Chua. 2022. Topic-Guided Conversational Recommender in Multiple Domains. *TKDE* (2022), 2485–2496.

[27] Lizi Liao, Tongyao Zhu, Le Hong Long, and Tat Seng Chua. 2021. Multi-domain dialogue state tracking with recursive inference. In *WWW*. 2568–2577.

[28] Ting-En Lin, Hua Xu, and Hanlei Zhang. 2020. Discovering new intents via constrained deep adaptive clustering with cluster refinement. 8360–8367.

[29] Jiqun Liu and Chirag Shah. 2019. Proactive identification of query failure. *ASIST* (2019), 176–185.

[30] Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. In *ACL*. 1036–1049.

[31] John P. Miller, Rohan Taori, Aditi Raghunathan, Shiori Sagawa, Pang Wei Koh, Vaishaal Shankar, Percy Liang, Yair Carmon, and Ludwig Schmidt. 2021. Accuracy on the Line: on the Strong Correlation Between Out-of-Distribution and In-Distribution Generalization. 7721–7735.

[32] Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. OpenDialKG: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. In *ACL*. 845–854.

[33] Helen Ngo, Cooper Raterink, João G. M. Araújo, Ivan Zhang, Carol Chen, Adrien Morisot, and Nicholas Frosst. 2021. Mitigating harm in language models with conditional-likelihood filtration. *CoRR* abs/2108.07790 (2021).

[34] Cennet Oguz and Ngoc Thang Vu. 2021. Few-shot Learning for Slot Tagging with Attentive Relational Network. In *EACL*. 1566–1572.

[35] Jinghui Qin, Zheng Ye, Jianheng Tang, and Xiaodan Liang. 2020. Dynamic knowledge routing network for target-guided open-domain conversation. In *AAAI*, Vol. 34. 8657–8664.

[36] Xuhui Ren, Hongzhi Yin, Tong Chen, Hao Wang, Zi Huang, and Kai Zheng. 2021. Learning to ask appropriate questions in conversational recommendation. In *SIGIR*. 808–817.

[37] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *preprint arXiv:1707.06347* (2017).

[38] Chirag Shah and Ryen W White. 2021. Applying Task Information for Search and Recommendations. In *Task Intelligence for Search and Recommendation*. 63–74.

[39] Chirag Shah and Ryen W White. 2021. Bridging task expressions and search queries. In *CHIIR*. 319–323.

[40] Darsh Shah, Raghav Gupta, Amir Fayazi, and Dilek Hakkani-Tur. 2019. Robust Zero-Shot Cross-Domain Slot Filling with Example Values. In *ACL*. 5484–5490.

[41] AB Siddique, Fuad Jamour, and Vagelis Hristidis. 2021. Linguistically-enriched and context-aware zero-shot slot filling. In *WWW*. 3279–3290.

[42] Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric Xing, and Zhiting Hu. 2019. Target-Guided Open-Domain Conversation. In *ACL*.

[43] Zhiwen Tang, Hrishikesh Kulkarni, and Grace Hui Yang. 2021. High-Quality Dialogue Diversification by Intermittent Short Extension Ensembles. In *Findings of ACL-IJCNLP*. 1861–1872.

[44] Zhiwen Tang and Grace Hui Yang. 2022. A Re-Classification of Information Seeking Tasks and Their Computational Solutions. *ACM Trans. Inf. Syst.* 40, 4, Article 80 (2022).

[45] Naftali Tishby, Fernando C. Pereira, and William Bialek. 1999. The information bottleneck method. In *Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing*. 368–377.

[46] Siddharth Verma, Justin Fu, Mengjiao Yang, and Sergey Levine. 2022. CHAI: A CHatbot AI for Task-Oriented Dialogue with Offline Reinforcement Learning.

[47] Marilyn Walker and Steve Whittaker. 1990. Mixed initiative in dialogue: an investigation into discourse segmentation. In *ACL*. 70–78.

[48] Dazhen Wan, Zheng Zhang, Qi Zhu, Lizi Liao, and Minlie Huang. 2022. A Unified Dialogue User Simulator for Few-shot Data Augmentation. In *EMNLP Findings*. To appear.

[49] Liwen Wang, Xuefeng Li, Jiachi Liu, Keqing He, Yuanmeng Yan, and Weiran Xu. 2021. Bridge to Target Domain by Prototypical Contrastive Learning and Label Confusion: Re-explore Zero-Shot Learning for Slot Filling. In *EMNLP*. 9474–9480.

[50] Xuewei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for Good: Towards a Personalized Persuasive Dialogue System for Social Good. In *ACL*. 5635–5649.

[51] Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. Proactive Human-Machine Conversation with Explicit Conversation Goal. In *ACL*. 3794–3804.

[52] Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. Proactive Human-Machine Conversation with Explicit Conversation Goal. In *ACL*. 3794–3804.

[53] Yuxia Wu, Lizi Liao, Xuemin Qian, and Tat-Seng Chua. 2022. Semi-supervised New Slot Discovery with Incremental Clustering. In *EMNLP Findings*. To appear.

[54] Yuxia Wu, Lizi Liao, Gangyi Zhang, Wenqiang Lei, Guoshuai Zhao, Xueming Qian, and Tat-Seng Chua. 2022. State Graph Reasoning for Multimodal Conversational Recommendation. *TMM* (2022).

[55] Yanan Wu, Zhiyuan Zeng, Keqing He, Hong Xu, Yuanmeng Yan, Huixing Jiang, and Weiran Xu. 2021. Novel Slot Detection: A Benchmark for Discovering Unknown Slot Types in the Task-Oriented Dialogue System. In *ACL-IJCNLP*.

[56] Jun Xu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Conversational graph grounded policy learning for open-domain conversation generation. In *ACL*. 1835–1845.

[57] Chenchen Ye, Lizi Liao, Fuli Feng, Wei Ji, and Tat-Seng Chua. 2022. Structured and Natural Responses Co-Generation for Conversational Search. In *SIGIR*. 155–164.

[58] Chenchen Ye, Lizi Liao, Suyu Liu, and Tat-Seng Chua. 2022. Reflecting on Experiences for Response Generation. In *ACM MM*.

[59] Zengfeng Zeng, Dan Ma, Haiqin Yang, Zhen Gou, and Jianping Shen. 2021. Automatic intent-slot induction for dialogue systems. In *WWW*. 2578–2589.

[60] B. Zhang, Z. Tu, Y. Jiang, S. He, G. Chao, D. Chu, and X. Xu. 2021. DGPF:A Dialogue Goal Planning Framework for Cognitive Service Conversational Bot. In *ICWS*. 335–340.

[61] Hanlei Zhang, Hua Xu, Ting-En Lin, and Rui Lyu. 2021. Discovering new intents with deep aligned clustering. In *AAAI*. 14365–14373.

[62] Jun Zhang, Yan Yang, Chencai Chen, Liang He, and Zhou Yu. 2021. KERS: A Knowledge-Enhanced Framework for Recommendation Dialog Systems with Multiple Subgoals. In *Findings of EMNLP*. 1092–1101.

[63] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing Dialogue Agents: I have a dog, do you have pets too?. In *ACL*. 2204–2213.

[64] Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020. Towards Topic-Guided Conversational Recommender System. In *COLING*. 4128–4139.

[65] Yutao Zhu, Jian-Yun Nie, Kun Zhou, Pan Du, Hao Jiang, and Zhicheng Dou. 2021. Proactive Retrieval-Based Chatbots Based on Relevant Knowledge and Goals. In *SIGIR*. 2000–2004.

[66] Jie Zou, Yifan Chen, and Evangelos Kanoulas. 2020. Towards question-based recommender systems. In *SIGIR*. 881–890.

[67] Jie Zou and Evangelos Kanoulas. 2019. Learning to ask: Question-based sequential Bayesian product search. In *CIKM*. 369–378.