

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

7-2023

Proactive conversational agents in the post-ChatGPT world

Lizi LIAO

Singapore Management University, lzliao@smu.edu.sg

Grace Hui YANG

Georgetown University

Chirag SHAH

University of Washington

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Higher Education Commons](#)

Citation

LIAO, Lizi; YANG, Grace Hui; and SHAH, Chirag. Proactive conversational agents in the post-ChatGPT world. (2023). *SIGIR '23: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, Taipei, July 23-27*. 3452-3455.

Available at: https://ink.library.smu.edu.sg/sis_research/8146

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.



Proactive Conversational Agents in the Post-ChatGPT World

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

ChatGPT and similar large language model (LLM) based conversational agents have brought shock waves to the research world. Although astonished by their human-like performance, we find they share a significant weakness with many other existing conversational agents in that they all take a passive approach in responding to user queries. This limits their capacity to understand the users and the task better and to offer recommendations based on a broader context than a given conversation. Proactiveness is still missing in these agents, including their ability to initiate a conversation, shift topics, or offer recommendations that take into account a more extensive context. To address this limitation, this tutorial reviews methods for equipping conversational agents with proactive interaction abilities.

The full-day tutorial is divided into four parts, including multiple interactive exercises. We will begin the tutorial with an interactive exercise and cover the design of existing conversational systems architecture and challenges. The content includes coverage of LLM-based recent advancements such as ChatGPT and Bard, along with reinforcement learning with human feedback (RLHF) technique. Then we will introduce the concept of proactive conversational agents and preset recent advancements in proactiveness of conversational agents, including actively driving conversations by asking questions, topic shifting, and methods that support strategic planning of conversation. Next, we will discuss important issues in conversational responses' quality control, including safety, appropriateness, language detoxication, hallucination, and alignment. Lastly, we will launch another interactive exercise and discussion with the audience to arrive at concluding remarks, prospecting open challenges and new directions. By exploring new techniques for enhancing conversational agents' proactive behavior to improve user engagement, this tutorial aims to help researchers and practitioners develop more effective conversational agents that can better understand and respond to user needs proactively and safely.

CCS CONCEPTS

• **Computing methodologies** → **Discourse, dialogue and pragmatics**; • **Information systems** → *Users and interactive retrieval*.

KEYWORDS

Proactive conversation, conversational search, conversational AI

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.

SIGIR '23, July 23–27, 2023, Taipei, Taiwan

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9408-6/23/07...\$15.00

<https://doi.org/10.1145/3539618.3594250>

ACM Reference Format:

Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. Proactive Conversational Agents in the Post-ChatGPT World. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '23)*, July 23–27, 2023, Taipei, Taiwan. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3539618.3594250>

1 MOTIVATION

The development of conversational agents that can comprehend human language and provide appropriate responses has long been a desired goal of Artificial Intelligence (AI). These agents can be broadly classified into two categories: (1) Chit-chat systems, which aim to engage users and offer emotional support by engaging in open-ended discussions on various topics, and (2) Task-oriented dialogue systems, which assist users in accomplishing specific tasks. Many commercial personal assistants, including Amazon Alexa, Apple Siri, Google Home, and large language model (LLM)-enabled Microsoft Copilot, fall under the task-oriented category. These systems are primarily designed to comprehend natural language verbal commands, interpret them, and translate them into actions to be executed by underlying application systems.

Recently, ChatGPT [1] and similar LLM-based conversational agents have brought shock waves to the research community and to the world. Astonished by their human-level performances, we notice that they share a significant weakness with most other existing conversational agents in that they all take a passive approach in responding to user queries. Their main research efforts are still on performing pre-defined actions or providing factual information in response to user commands or queries. This limits their capacity to understand the users and the task better and to offer recommendations based on a broader context than a given conversation. The missing proactiveness includes lacking abilities to initiate a conversation, shift topics, strategic plan with subgoals, or offer recommendations that take into account a more extensive context beyond the scope of a specific conversation.

Moreover, despite being widely adopted and receiving tremendous attention, most current conversational agents, including LLM-enabled ones, heavily rely on pre-existing training conversations, datasets, and knowledge associated with them in order to exchange information, provide recommendations, and complete tasks [1, 24, 26, 32, 53]. They typically generate responses to questions in a passive manner rather than leading the conversation or asking questions themselves [9, 45, 46]. This reactive, passive approach to conversation limits the range of conversations that can take place, particularly in situations that require active engagement from both sides, such as exploratory search or complex decision-making.

In recent years, researchers from multiple fields, including natural language processing, dialogue systems, and machine learning [12, 16, 19, 33, 52, 54], have been working towards the goal of enabling conversational agents to engage in two-way, proactive

conversations with users [28, 30, 47, 51, 55]. They have proposed various approaches to address this issue, such as:

- Learning to ask [5, 13, 22, 25, 36, 44, 50, 56, 57]
- Topic shifting [27, 39, 49]
- Strategy planning with reinforcement Learning, counterfactual dialogue act, and label generation [2, 8, 22, 29, 40, 42].

In addition to response accuracy and effectiveness, other important aspects about conversational agents have also started to attract research attention in recent years. These aspects are particularly relevant to proactive conversational agents because these agents are taking more control over the conversation. They span topics on

- Conversational safety
- Response appropriateness
- Language detoxication
- Language model's hallucination; and
- AI's alignment issues.

These matters can be viewed as under the umbrella of conversational responses quality control [3, 4, 15, 18, 31, 37, 40, 43].

Furthermore, there have been efforts to improve evaluation methods for conversational agents [21, 41].

In this tutorial, we provide a comprehensive review of the literature on this emerging research topic and discuss structured methods to equip conversational agents with the ability to proactively interact with users to perform various tasks. These methods are aimed at overcoming the limitations of current conversational agents, which can only respond reactively and passively to user commands and queries. By enhancing conversational agents with the ability to engage in proactive conversations, we can improve their utility and effectiveness in a range of settings, including exploratory search and complex decision-making.

2 OBJECTIVES

Content-wise, this tutorial aims to

- Offer a thorough review of the literature on this developing research area and explore structured techniques for enabling conversational agents to actively interact with users and execute various tasks;
- Address the constraints of existing conversational agents that solely respond reactively and passively to user instructions and inquiries and enhance conversational agents with proactive conversation capabilities;
- Increase the proactiveness of conversational agents in diverse contexts, such as exploratory search and complex decision-making processes

This full-day tutorial will be divided into four parts, including 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 how to drive conversation by asking questions actively and purposefully making topic shifts. In the third part, we review and discuss methods that support strategic planning by reinforcement learning-powered methods and

retrospectively conducting response quality control. Important issues, such as hallucination, language detoxication, and the agents' potential to be power-seeking, arising from the popularity of conversational agents and especially LLM-based conversational agents will be reviewed and discussed. The latest proposed solutions to resolve these newly-arising issues will be the focus of this part of the tutorial. 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 with a discussion on the scopes for proactive conversational agents and present some open challenges for future research.

Teaching-wise, the tutorial aims to provide vivid presentations and hands-on experiences, which will include the use of conversational devices such as Amazon Alexa and Apple Siri in some of our examples and interactive exercises with attendees. These devices will be installed in the tutorial room to demonstrate the capabilities of conversational agents and to provide attendees with an immersive learning experience. In addition, an online discussion forum will be created concurrently with the tutorial. Attendees' questions will be collected during and after the tutorial. We will post answers to the questions and maintain the forum during the entire conference.

3 RELEVANCE TO THE IR COMMUNITY

In recent years, there has been a growing interest in conversational agents within the fields of information retrieval (IR), artificial intelligence (AI), and natural language processing (NLP). This is evidenced by the inclusion of a Conversational Assistant Track in the TREC conferences from 2019 to 2022, as well as the recognition of "Dialogue and Interactive Systems" as a major research topic at SIGIR, a leading IR conference.

Our tutorial aims to address the pressing need for post-ChatGPT conversational agents that can engage in proactive conversations and actively contribute to ongoing interactions. Given the widespread use of conversational agents in customer service, task completion, and personal assistant applications, there is an urgent demand to enhance their proactivity. As conversational interfaces continue to gain traction in both commercial and personal contexts, we believe that our tutorial will be of significant interest and benefit to both IR researchers and practitioners.

Similar Tutorials in Related Conferences:

In recent related conferences, several tutorials on general dialogue systems and conversational recommendation have been given, for example, *Conversational Information Seeking* [11] in SIGIR'22 and *Conversational Recommendation Systems* [17] in RecSys'20 and in WSDM'21. However, these tutorials have not focused specifically on the proactive aspect of conversational agents. Our tutorial provides a new perspective on conversational agents and expands the current understanding of this field.

Our tutorial covers topics that are complementary to existing mainstream approaches, including topic shifting, inappropriate dialogue management, and conservative reinforcement learning. To the best of our knowledge, our tutorial is the first to specifically focus on proactive conversational agents in SIGIR. This tutorial is

particularly timely as it addresses an urgent need to improve conversational agents in commercial and daily use cases. The content of this tutorial is expected to be of interest to both researchers and practitioners in the field of information retrieval, artificial intelligence, and natural language processing.

4 FORMAT AND DETAILED SCHEDULE

The tutorial will comprise primarily of didactic lectures, complemented by interactive exercises designed for both in-person and remote attendees. Additionally, we will facilitate a virtual forum to facilitate interactive question-and-answer sessions and foster discussion among participants both during and following the tutorial.

Prior to the conference, we will engage in promotional activities for our tutorial, including leveraging social media channels and professional email lists such as the SIGIR mailing list. We will also create and maintain a dedicated tutorial website to promote the content and share news with colleagues across various research areas, including NLP, AI, Speech, Dialogue Systems, ML, Recommender systems, and IR.

During the tutorial, we plan to include multiple interactive exercises with attendees to enhance learning and engagement. To facilitate communication and knowledge transfer, we will actively encourage attendees to submit questions and participate in a forum that we will create specifically for the tutorial, which will be maintained throughout the conference.

In the first interactive exercise 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 Taipei. 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 10 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 Four, 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.

ACKNOWLEDGEMENTS

This research was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant.

Table 1: Schedule of the Tutorial [6 hrs].

Part One [90 min]
1) Introduction
1.1) Motivation & Aim of this tutorial
2) Interactive Exercise [20 min]
3) Proactive Conv. Agents: Overview
3.1) Existing conversational agents
3.1.1) Non-LLM conversational agents
3.1.2) LLM conversational agents [30 min] [6, 32, 34, 35]
3.2) Applications: Search, recommend, and more [20, 48]
3.3) Evaluation: Metrics, procedures, and challenges [21, 41]
<i>Coffee break</i>
Part Two [90 min]
4) Learning to Ask
4.1) Mixed initiatives [22, 44]
4.2) Learning to ask [25, 36, 56, 57]
4.3) Question selection and generation [5, 13, 50]
5) Topic Shifting
5.1) Target-guided open chitchat
5.2) Target-guided conversational recommendation [27, 39, 49]
<i>Lunch break</i>
Part Three [90 min]
6) Strategy Planning
6.1) Reinforcement Learning with subgoals
6.2) Counterfactual dialogue act augmentation [8, 22, 40, 42]
6.3) Out-of-distribution label generation [2, 29]
7) Response Quality Control
7.1) Types of inappropriate responses [14]
7.2) Language detoxification [3, 15, 31]
7.3) Conservative reinforcement learning [8, 10, 23, 38, 40, 43]
7.4) Combat power-seeking AI [7]
<i>Coffee break</i>
Part Four: [90 min]
8) Interactive Exercise-2 [20 min]
9) Conclusion
9.1) Interactive & Future directions
9.2) Challenges & Future directions

REFERENCES

- [1] Open AI. 2022. ChatGPT. <https://openai.com/blog/chatgpt/>
- [2] Christina Baek, Yiding Jiang, Aditi Raghunathan, and Zico Kolter. 2022. Agreement-on-the-Line: Predicting the Performance of Neural Networks under Distribution Shift.
- [3] 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.
- [4] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson

- Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauha Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. [arXiv:2204.05862](https://arxiv.org/abs/2204.05862) [cs.CL]
- [5] Keping Bi, Qingyao Ai, and W. Bruce Croft. 2021. Asking Clarifying Questions Based on Negative Feedback in Conversational Search. In *SIGIR '21*. ACM.
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. 1877–1901.
- [7] Joseph Carlsmith. 2022. Is Power-Seeking AI an Existential Risk? [arXiv:2206.13353](https://arxiv.org/abs/2206.13353) [cs.CY]
- [8] Limin Chen, Zhiwen Tang, and Grace Hui Yang. 2020. Balancing Reinforcement Learning Training Experiences in Interactive Information Retrieval. In *SIGIR*. 1525–1528.
- [9] Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. 2022. SalesBot: Transitioning from Chat-Chat to Task-Oriented Dialogues. In *ACL*. 6143–6158.
- [10] 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.
- [11] 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*. 3455–3458.
- [12] Jeffrey Dalton, Chenyan Xiong, and Jamie Callan. 2021. TREC CAsT 2021: The Conversational Assistance Track Overview. In *TREC*.
- [13] 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.
- [14] 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.
- [15] 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*.
- [16] 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.
- [17] Zuohui Fu, Yikun Xian, Yongfeng Zhang, and Yi Zhang. 2020. Tutorial on Conversational Recommendation Systems. In *RecSys*.
- [18] Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y. Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2022. RARR: Researching and Revising What Language Models Say, Using Language Models. [arXiv:2210.08726](https://arxiv.org/abs/2210.08726) [cs.CL]
- [19] Prakhhar Gupta, Jeffrey P Bigham, Yulia Tsvetkov, and Amy Pavel. 2021. Controlling Dialogue Generation with Semantic Exemplars. In *NAACL*. 3018–3029.
- [20] 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.
- [21] Chathra Hendahewa and Chirag Shah. 2017. Evaluating user search trails in exploratory search tasks. *IP&M* (2017), 905–922.
- [22] 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*.
- [23] Aviral Kumar, Aurick Zhou, George Tucker, and Sergey Levine. 2020. Conservative Q-Learning for Offline Reinforcement Learning. *CoRR abs/2006.04779* (2020).
- [24] 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.).
- [25] Zixuan Li, Lizi Liao, and Tat-Seng Chua. 2022. Learning to Ask Critical Questions for Assisting Product Search. In *SIGIR Workshop on eComm*.
- [26] Lizi Liao, Yunshan Ma, Xiangnan He, Richang Hong, and Tat-seng Chua. 2018. Knowledge-aware multimodal dialogue systems. In *ACM MM*. 801–809.
- [27] Lizi Liao, Ryuichi Takanobu, Yunshan Ma, Xun Yang, Minlie Huang, and Tat-Seng Chua. 2022. Topic-Guided Conversational Recommender in Multiple Domains. *TKDE* (2022), 2485–2496.
- [28] 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.
- [29] 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.
- [30] Seungwhan Moon, Pararth Shah, Anuj Kumar, and Rajen Subba. 2019. OpenDialog: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. In *ACL*. 845–854.
- [31] Helen Ngo, Cooper Raterink, João G. M. Araújo, Ivan Zhang, Carol Chen, Adrien Moriso, and Nicholas Frosst. 2021. Mitigating harm in language models with conditional-likelihood filtration. *CoRR abs/2108.07790* (2021).
- [32] OpenAI. 2023. GPT-4 Technical Report. [arXiv:2303.08774](https://arxiv.org/abs/2303.08774) [cs.CL]
- [33] 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.
- [34] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. (2018).
- [35] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. Language Models are Unsupervised Multitask Learners. (2018).
- [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] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. Toolformer: Language models can teach themselves to use tools. *arXiv preprint arXiv:2302.04761* (2023).
- [38] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *preprint arXiv:1707.06347* (2017).
- [39] Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric Xing, and Zhiting Hu. 2019. Target-Guided Open-Domain Conversation. In *ACL*.
- [40] 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.
- [41] 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).
- [42] 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*.
- [43] Siddharth Verma, Justin Fu, Mengjiao Yang, and Sergey Levine. 2022. CHAI: A Chatbot AI for Task-Oriented Dialogue with Offline Reinforcement Learning.
- [44] Marilyn Walker and Steve Whittaker. 1990. Mixed initiative in dialogue: an investigation into discourse segmentation. In *ACL*. 70–78.
- [45] 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.
- [46] 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.
- [47] 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.
- [48] 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).
- [49] 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.
- [50] 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.
- [51] B. Zhang, Z. Tu, Y. Jiang, S. He, G. Chao, D. Chu, and X. Xu. 2021. DGPFA: Dialogue Goal Planning Framework for Cognitive Service Conversational Bot. In *ICWS*. 335–340.
- [52] 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.
- [53] 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.
- [54] Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020. Towards Topic-Guided Conversational Recommender System. In *COLING*. 4128–4139.
- [55] 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.
- [56] Jie Zou, Yifan Chen, and Evangelos Kanoulas. 2020. Towards question-based recommender systems. In *SIGIR*. 881–890.
- [57] Jie Zou and Evangelos Kanoulas. 2019. Learning to ask: Question-based sequential Bayesian product search. In *CIKM*. 369–378.