11-2023

Rethinking conversational agents in the era of large language models: Proactivity, non-collaborativity, and beyond

Yang DENG
Singapore Management University, ydeng@smu.edu.sg

Wenqiang LEI

Minlie HUANG

Tat-Seng CHUA

Rethinking Conversational Agents in the Era of Large Language Models: Proactivity, Non-collaborativity, and Beyond

Yang Deng
National University of Singapore
Singapore
ydeng@nus.edu.sg

Wenqiang Lei
Sichuan University
Chengdu, China
wenqianglei@gmail.com

Minlie Huang
Tsinghua University
Beijing, China
aihuang@tsinghua.edu.cn

Tat-Seng Chua
National University of Singapore
Singapore
chuats@comp.nus.edu.sg

ABSTRACT

Conversational systems are designed to offer human users social support or functional services through natural language interactions. Typical conversation researches mainly focus on the responseability of the system, such as dialogue context understanding and response generation. In the era of large language models (LLMs), LLM-augmented conversational systems showcase exceptional capabilities of responding to user queries for different language tasks. However, as LLMs are trained to follow users' instructions, LLM-augmented conversational systems typically overlook the design of an essential property in intelligent conversations, *i.e.*, goal awareness. In this tutorial, we will introduce the recent advances on the design of agent's awareness of goals in a wide range of conversational systems, including proactive, non-collaborative, and multi-goal conversational systems. In addition, we will discuss the main open challenges in developing agent's goal awareness in LLM-augmented conversational systems and several potential research directions for future studies.

CCS CONCEPTS

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

KEYWORDS

Open-domain Dialogue, Task-oriented Dialogue, Conversational Information Seeking, Proactivity

ACM Reference Format:
Yang Deng, Wenqiang Lei, Minlie Huang, and Tat-Seng Chua. 2023. Rethinking Conversational Agents in the Era of Large Language Models: Proactivity, Non-collaborativity, and Beyond. In *Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region (SIGIR-AP '23), November 26–28, 2023, Beijing, China.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3624918.3629548

SIGIR-AP '23, November 26–28, 2023, Beijing, China

ACM ISBN 979-8-4007-0408-6/23/11…$15.00
https://doi.org/10.1145/3624918.3629548

1 MOTIVATION AND OVERVIEW

Conversational systems are envisioned to provide social support or functional service to human users via natural language interactions. Conversation researches typically center around a system's response capabilities, such as understanding the context of dialogue and generating appropriate responses to user requests. The popularity of conversational systems has grown unprecedentedly with the advent of ChatGPT [39], which showcases exceptional proficiency in the capabilities of context understanding and response generation with large language models (LLMs). Meanwhile, its powerfulness also raises questions about the potential for advancement to take the conversational system to the next level.

However, existing conversational systems often neglect to emphasize an integral property in intelligent conversations: *goal awareness*. The goal awareness means the state of not only being responsive to the users but also aware of the target conversational goal and capable of leading the conversation towards the goal, which is a significant step towards strong artificial intelligence. It can not only largely improve user engagement and service efficiency in the conversation, but also empower the system to handle more complicated conversation tasks that involve strategic and motivational interactions. The main content of this tutorial includes:

1.1 Proactive Conversational Systems

As opposed to responding to users, proactivity is the most prominent feature of goal awareness in conversational systems, which can improve the collaboration between the users and system towards the ultimate conversation goal. Derived from the definition of proactivity in organizational behaviors [23] and its dictionary definitions, conversational agents' proactivity can be defined as the capability to create or control the conversation by taking the initiative and anticipating impacts on themselves or human users. In this part, we will provide a comprehensive introduction about such efforts on the design of agent's proactivity that span various task formulations and application scenarios. In specific, we categorize them in three directions according to the application scenario, and plan to discuss their research problems and methods as follows:

- **Topic Shifting and Planning in Open-domain Dialogues**. The goal of ODD systems is to maintain engaging social conversations with users. Proactive ODD systems can consciously change topics [49] and lead directions [45, 48] for improving

user engagement in the conversation. We will present the existing methods for topic shifting and planning in open-domain dialogues, including keyword-based discourse-level topic planning [45], graph-based topic planning [38, 52], and learning from interactions with users [28].

- **Additional Information Delivery in Task-oriented Dialogues**. The goal of TOD systems is to provide functional service for users, such as making reservations or managing schedule. The proactivity in TOD systems is firstly defined as the capability of consciously providing additional information that is not requested by but useful to the users [3], which can improve the quality and effectiveness of conveying functional service in the conversation. We will introduce the recent studies of proactive TOD systems with various designs, including adding topical chitchats into the responses for more engaging interactions [44] and enriching TOD with relevant entity knowledge [8].
- **Uncertainty Elimination in Information-seeking Dialogues**. The goal of CIS systems [58] is to fulfill the user's information needs and its typical applications include conversational search, conversational recommendation, and conversational question answering. Conventional CIS systems assume that users always convey clear information requests, while the user queries, in reality, are often brief and succinct. Recent years have witnessed several advances on developing proactive CIS systems that can consciously eliminate the uncertainty for more efficient and precise information seeks by initiating a subdialogue. Such a subdialogue can either clarify the ambiguity of the query or question in conversational search [1, 57] and conversation question answering [15, 24], or elicit the user preference in conversational recommendation [27, 62].

1.2 Non-collaborative Conversational Systems

Most of existing conversational systems are built upon the assumption that the users willingly collaborate with the conversational agent to reach the mutual goal. However, this assumption may not always hold in some real-world scenarios, where the users and the system do not share the same goal [25, 47] or the users are not willing to coordinate with the agent [26, 51]. In these cases, the conversational agent requires another feature of goal awareness, *i.e.*, non-collaborativity [30, 64], which means the capability of handling both in-goal and off-goal dialogues appropriately for ultimately leading back to the system's goal. In this part, we will categorize the non-collaborative settings into two groups as follows and cover their to-date work respectively.

- **The users are not willing to coordinate with the agent.** Example scenarios include calming down the emotional users before solving their problems [18, 34], managing the users' complaints before providing service [51], and handling problematic content during the conversations [26]. We will introduce the pioneering studies for the system to consciously deal with non-collaborative users during the conversation, including emotion cause analysis [9, 46], user satisfaction estimation [33], and prosocial response generation [2, 26].
- **The users and the system do not share the same goal.** Typical applications include persuasion dialogues [47], negotiation dialogues [25], and anti-scam dialogues [30]. We will present the approaches for the system to consciously mitigate and resolve the conflict goals with users, including dialogue strategy learning [30, 64], user personality modeling [50], and persuasive response generation [37].

1.3 Multi-goal Conversational Systems

All the aforementioned conversational systems assume that users always know what they want and the system solely targets at reaching a certain goal, such as chit-chat, question answering, recommendation, etc. The system with a higher level of agent's awareness of goals should also be capable of handling conversations with multiple and various goals. As for multi-goal conversational systems [17, 36], the agent is expected to consciously discover users' intentions and naturally lead user-engaged dialogues with multiple conversation goals. We will cover the newly proposed problems in multi-goal conversational systems with their corresponding data resources [10, 35, 55]. Then we will discuss two problem settings of multi-goal conversational systems with corresponding state-of-the-art approaches: (i) The goal sequence is pre-defined [59], and (ii) The next goal needs to be predicted [17, 35].

1.4 Open Challenges and Beyond

In the last part, we will discuss the main open challenges in developing agent's awareness in conversational systems and several potential research directions for future studies.

- **Goal Awareness in LLM-augmented Conversational AI.** Large Language Models (LLMs) have been demonstrated to be powerful of handling various NLP tasks in the form of conversations, such as ChatGPT. However, these applications are typically designed to follow the user's instructions and conversational intents. There are still several limitations that attribute to the lack of agent's awareness, such as passively providing randomly-guessed answers to ambiguous user queries, failing to refuse or handle problematic user requests that may exhibit harmful or biased conversations, etc. In addition, they also fall short of interacting under non-collaborative or system-oriented settings. Therefore, we will discuss the latest studies in triggering the proactivity of LLM-based dialogue systems [6, 14, 19, 56, 60, 63] and the planning capabilities of LLMs with some latest studies [41, 53, 54].
- **Evaluation for Conversational Agent's Goal Awareness.** The development of robust evaluation protocols has already been a long-standing problem for different kinds of conversational systems. The evaluation for conversational agent's awareness is a more challenging problem, since it is involved the evaluation not only from the perspective of natural language, but also from the perspectives of human-computer interaction, sociology, psychology, etc. We will cover the latest studies for shedding some lights on this topic, inclusive of popular metrics such as goal completion and user satisfaction [16, 35, 52], and model-based methods such as user simulator [40, 61].
- **Ethics for Conversational Agent's Goal Awareness.** Although existing designs of agent's goal awareness in conversational systems generally aim at social goodness [26, 47], it is inevitably a double-edged sword that can be used for good or evil. For responsible researches, we will discuss several important aspects of ethical issues in conversational systems: (i) Factuality: Factual

incorrectness and hallucination are common in conversational systems [4]. When enabling the conversational agent with awareness, it becomes more crucial to guarantee the factuality of the system-provided information [5]. (ii) Safety: Besides general dialogue safety problems, such as toxic language and social bias [43], conscious conversational systems need to pay more attentions to the aggressiveness issue during the non-collaborative conversations [25, 34]. (iii) Privacy: The privacy issue is overlooked in current studies on conversational systems [29], but the agent's awareness raises concerns about how these conversational systems handle personal information obtained from the users.

2 OBJECTIVES

The main objectives of this tutorial are threefold:

- This tutorial presents a comprehensive and diverse overview about the cutting-edge designs of agent's awareness in various conversational systems. The discussed approaches are problem-driven and language-agnostic, which means that the techniques are also not limited to a certain type of dialogues and can be generalized to diverse conversational systems.
- This tutorial discusses open challenges for goal awareness in conversational AI. LLMs have showcased exceptional proficiency in enhancing the response-ability of conversational systems. We provide a new perspective to facilitate more potential directions for future research into conversational AI.
- This tutorial provides the opportunity to arouse discussions on conversational AI's awareness of goals from the view of ethical and responsible researches. As part of this tutorial, we will provide a specific section for discussing the ethical considerations and designs for agent's awareness in conversational systems.

3 FORMAT AND DETAILED SCHEDULE

The following summarizes the detailed schedule of the tutorial:

(1) Introduction [10 min]
(2) Conversational System Preliminaries [20 min]
  (a) Open-domain Dialogue Systems
  (b) Task-oriented Dialogue Systems
  (c) Conversational Information-seeking Systems
(3) Proactive Conversational Systems [60 min]
  (a) Topic Shifting and Planning in Open-domain Dialogues
  (b) Additional Information Delivery in Task-oriented Dialogues
  (c) Uncertainty Elimination in Information-seeking Dialogues
(4) Non-collaborative Conversational Systems [40 min]
  (a) The users are not willing to coordinate with the system
  (b) The users and the system do not share the same goal
(5) Multi-goal Conversational Systems [20 min]
(6) Open Challenges for Conversational Agents' Goal Awareness and Beyond [20 min]
  (a) Goal Awareness in LLM-augmented Conversational AI
  (b) Evaluation for Conversational Agent's Goal Awareness
  (c) Ethics for Conversational Agent's Goal Awareness
(7) Summary and Outlook [10 min]

4 RELEVANCE TO IR COMMUNITY

The conversational system is an trending topic in the information retrieval community, which receives notably increasing attentions from both academia and industry. In academia, IR conferences recognize *Conversational Systems* as one of its major research topics and host regular sessions about conversational systems. In industry, recent years have witnessed many successful applications that evolve traditional interactive IR systems into conversational IR systems. For example, Microsoft recently released a new version of Bing with its integration with ChatGPT [39] under the idea of conversational search. Our tutorial aims at stimulating progresses on conversational systems to the next level by jumping out of the box of reactive conversational systems that simply respond to the user requests. We focus on an important feature in conversational systems towards higher-level intelligence and artificial consciousness, which is the *goal awareness*.

Several tutorials about dialogue systems [7, 11, 20–22, 42] in general have been given in various conferences. However, these tutorials mainly introduce the advanced designs on the response-ability of the conversational system for general dialogue problem settings. The *Conversational Information Seeking: Theory and Application* tutorial [11] includes a section about mixed-initiative interactions in conversational information-seeking systems to present the recent studies on asking clarification questions in conversational search. The *Proactive Conversational Agents* tutorial [31, 32] was presented at WSDM/SIGIR 2023, which is the most relevant tutorial to our topics. In our tutorial, we identify the proactivity as a prominent feature of goal awareness in conversational systems and provide more comprehensive perspectives on the emerging problems and approaches that rely on higher level of conversational agents' awareness of goals, including proactive, non-collaborative, and multi-goal conversational systems. A previous version [12] of this tutorial was presented in ACL 2023, which is further incorporated with more content about LLM-based dialogue systems.

5 SUPPORTING MATERIALS

(1) **Slides** will be made publicly available; (2) The tutorial is accompanied with **a survey** [13] on this topic.

REFERENCES

[1] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. 2019. Asking Clarifying Questions in Open-Domain Information-Seeking Conversations. In *SIGIR 2019*. 475–484.
[2] Ashutosh Baheti, Maarten Sap, Alan Ritter, and Mark O. Riedl. 2021. Just Say No: Analyzing the Stance of Neural Dialogue Generation in Offensive Contexts. In *EMNLP 2021*. 4846–4862.
[3] Vevake Balaraman and Bernardo Magnini. 2020. Proactive systems and influenceable users: Simulating proactivity in task-oriented dialogues. In *SEMDIAL*. 1–12.
[4] Liang Chen, Yang Deng, Yatao Bian, Zeyu Qin, Bingzhe Wu, Tat-Seng Chua, and Kam-Fai Wong. 2023. Beyond Factuality: A Comprehensive Evaluation of Large Language Models as Knowledge Generators. *arXiv preprint arXiv:2310.07289* (2023).
[5] Maximillian Chen, Weiyan Shi, Feifan Yan, Ryan Hou, Jingwen Zhang, Saurav Sahay, and Zhou Yu. 2022. Seamlessly Integrating Factual Information and Social Content with Persuasive Dialogue. In *AACL/IJCNLP 2022*. 399–413.
[6] Maximillian Chen, Xiao Yu, Weiyan Shi, Urvi Awasthi, and Zhou Yu. 2023. Controllable Mixed-Initiative Dialogue Generation through Prompting. In *ACL 2023*. 951–966.
[7] Yun-Nung Chen, Asli Celikyilmaz, and Dilek Hakkani-Tür. 2017. Deep Learning for Dialogue Systems. In *ACL 2017: Tutorial Abstracts*. 8–14.
[8] Zhiyu Chen, Bing Liu, Seungwhan Moon, Chinnadhurai Sankar, Paul A. Crook, and William Yang Wang. 2022. KETOD: Knowledge-Enriched Task-Oriented Dialogue. In *Findings of ACL: NAACL 2022*. 2581–2593.
[9] Yi Cheng, Wenge Liu, Wenjie Li, Jiashuo Wang, Ruihui Zhao, Bang Liu, Xiaodan Liang, and Yefeng Zheng. 2022. Improving Multi-turn Emotional Support

Dialogue Generation with Lookahead Strategy Planning. *CoRR* abs/2210.04242 (2022).
[10] Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. 2022. SalesBot: Transitioning from Chit-Chat to Task-Oriented Dialogues. In *ACL 2022*. 6143–6158.
[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 '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 3455–3458.
[12] Yang Deng, Wenqiang Lei, Minlie Huang, and Tat-Seng Chua. 2023. Goal Awareness for Conversational AI: Proactivity, Non-collaborativity, and Beyond. In *ACL 2023: Tutorial Abstracts*. 1–10.
[13] Yang Deng, Wenqiang Lei, Wai Lam, and Tat-Seng Chua. 2023. A Survey on Proactive Dialogue Systems: Problems, Methods, and Prospects. In *IJCAI 2023*. 6583–6591.
[14] Yang Deng, Wenqiang Lei, Lizi Liao, and Tat-Seng Chua. 2023. Prompting and Evaluating Large Language Models for Proactive Dialogues: Clarification, Target-guided, and Non-collaboration. *CoRR* abs/2305.13626 (2023).
[15] Yang Deng, Wenqiang Lei, Wenxuan Zhang, Wai Lam, and Tat-Seng Chua. 2022. PACIFIC: Towards Proactive Conversational Question Answering over Tabular and Textual Data in Finance. *CoRR* abs/2210.08817 (2022).
[16] Yang Deng, Wenxuan Zhang, Wai Lam, Hong Cheng, and Helen Meng. 2022. User Satisfaction Estimation with Sequential Dialogue Act Modeling in Goal-oriented Conversational Systems. In *WWW 2022*. 2998–3008.
[17] Yang Deng, Wenxuan Zhang, Weiwen Xu, Wenqiang Lei, Tat-Seng Chua, and Wai Lam. 2023. A Unified Multi-task Learning Framework for Multi-goal Conversational Recommender Systems. *ACM Trans. Inf. Syst.* 41, 3 (2023), 77:1–77:25.
[18] Yang Deng, Wenxuan Zhang, Yifei Yuan, and Wai Lam. 2023. Knowledge-enhanced Mixed-initiative Dialogue System for Emotional Support Conversations. In *ACL 2023*. 4079–4095.
[19] Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. Improving Language Model Negotiation with Self-Play and In-Context Learning from AI Feedback. *CoRR* abs/2305.10142 (2023).
[20] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *ACL 2018, Tutorial Abstracts*. 2–7.
[21] Jianfeng Gao, Michel Galley, and Lihong Li. 2018. Neural Approaches to Conversational AI. In *SIGIR 2018*. 1371–1374.
[22] Jianfeng Gao, Chenyan Xiong, and Paul Bennett. 2020. Recent Advances in Conversational Information Retrieval. In *SIGIR 2020*. 2421–2424.
[23] Adam M Grant and Susan J Ashford. 2008. The dynamics of proactivity at work. *Research in organizational behavior* 28 (2008), 3–34.
[24] Meiqi Guo, Mingda Zhang, Siva Reddy, and Malihe Alikhani. 2021. Abg-CoQA: Clarifying Ambiguity in Conversational Question Answering. In *AKBC 2021*.
[25] He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling Strategy and Generation in Negotiation Dialogues. In *EMNLP 2018*. 2333–2343.
[26] Hyunwoo Kim, Youngjae Yu, Liwei Jiang, Ximing Lu, Daniel Khashabi, Gunhee Kim, Yejin Choi, and Maarten Sap. 2022. ProsocialDialog: A Prosocial Backbone for Conversational Agents. *CoRR* abs/2205.12688 (2022).
[27] Wenqiang Lei, Xiangnan He, Yisong Miao, Qingyun Wu, Richang Hong, Min-Yen Kan, and Tat-Seng Chua. 2020. Estimation-Action-Reflection: Towards Deep Interaction Between Conversational and Recommender Systems. In *WSDM 2020*.
[28] Wenqiang Lei, Yao Zhang, Feifan Song, Hongru Liang, Jiaxin Mao, Jiancheng Lv, Zhenglu Yang, and Tat-Seng Chua. 2022. Interacting with Non-Cooperative User: A New Paradigm for Proactive Dialogue Policy. In *SIGIR 2022*. 212–222.
[29] Haoran Li, Yangqiu Song, and Lixin Fan. 2022. You Don't Know My Favorite Color: Preventing Dialogue Representations from Revealing Speakers' Private Personas. In *NAACL 2022*. 5858–5870.
[30] Yu Li, Kun Qian, Weiyan Shi, and Zhou Yu. 2020. End-to-End Trainable Non-Collaborative Dialog System. In *AAAI 2020*. 8293–8302.
[31] Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. Proactive Conversational Agents. In *WSDM 2023*. 1244–1247.
[32] Lizi Liao, Grace Hui Yang, and Chirag Shah. 2023. Proactive Conversational Agents in the Post-ChatGPT World. In *SIGIR 2023*. 3452–3455.
[33] Jiawei Liu, Kaisong Song, Yangyang Kang, Guoxiu He, Zhuoren Jiang, Changlong Sun, Wei Lu, and Xiaozhong Liu. 2021. A Role-Selected Sharing Network for Joint Machine-Human Chatting Handoff and Service Satisfaction Analysis. In *EMNLP 2021*. 9731–9741.
[34] Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. Towards Emotional Support Dialog Systems. In *ACL/IJCNLP 2021*. 3469–3483.
[35] Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020. Towards Conversational Recommendation over Multi-Type Dialogs. In *ACL 2020*. 1036–1049.
[36] Zeming Liu, Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, and Hua Wu. 2022. Where to Go for the Holidays: Towards Mixed-Type Dialogs for Clarification of User Goals. In *ACL 2022*. 1024–1034.
[37] Kshitij Mishra, Azlaan Mustafa Samad, Palak Totala, and Asif Ekbal. 2022. PEPDS: A Polite and Empathetic Persuasive Dialogue System for Charity Donation. In *COLING 2022*. 424–440.
[38] Jinjie Ni, Vlad Pandelea, Tom Young, Haicang Zhou, and Erik Cambria. 2022. HiTKG: Towards Goal-Oriented Conversations via Multi-Hierarchy Learning. In *AAAI 2022*. 11112–11120.
[39] J Schulman, B Zoph, C Kim, J Hilton, J Menick, J Weng, JFC Uribe, L Fedus, L Metz, M Pokorny, et al. 2022. ChatGPT: Optimizing language models for dialogue.
[40] Ivan Sekulic, Mohammad Aliannejadi, and Fabio Crestani. 2022. Evaluating Mixed-initiative Conversational Search Systems via User Simulation. In *WSDM 2022*. 888–896.
[41] Noah Shinn, Beck Labash, and Ashwin Gopinath. 2023. Reflexion: an autonomous agent with dynamic memory and self-reflection. *CoRR* abs/2303.11366 (2023).
[42] Pei-Hao Su, Nikola Mrkšić, Iñigo Casanueva, and Ivan Vulić. 2018. Deep Learning for Conversational AI. In *NAACL 2018: Tutorial Abstracts*. 27–32.
[43] Hao Sun, Guangxuan Xu, Jiawen Deng, Jiale Cheng, Chujie Zheng, Hao Zhou, Nanyun Peng, Xiaoyan Zhu, and Minlie Huang. 2022. On the Safety of Conversational Models: Taxonomy, Dataset, and Benchmark. In *Findings of ACL: ACL 2022*. 3906–3923.
[44] Kai Sun, Seungwhan Moon, Paul A. Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho, and Claire Cardie. 2021. Adding Chit-Chat to Enhance Task-Oriented Dialogues. In *NAACL-HLT 2021*. 1570–1583.
[45] Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. 2019. Target-Guided Open-Domain Conversation. In *ACL 2019*. 5624–5634.
[46] Quan Tu, Yanran Li, Jianwei Cui, Bin Wang, Ji-Rong Wen, and Rui Yan. 2022. MISC: A Mixed Strategy-Aware Model integrating COMET for Emotional Support Conversation. In *ACL 2022*. 308–319.
[47] 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 2019*. 5635–5649.
[48] 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 2019*. 3794–3804.
[49] Huiyuan Xie, Zhenghao Liu, Chenyan Xiong, Zhiyuan Liu, and Ann A. Copestake. 2021. TIAGE: A Benchmark for Topic-Shift Aware Dialog Modeling. In *Findings of ACL: EMNLP 2021*. 1684–1690.
[50] Runzhe Yang, Jingxiao Chen, and Karthik Narasimhan. 2021. Improving Dialog Systems for Negotiation with Personality Modeling. In *ACL/IJCNLP 2021*.
[51] Wei Yang, Luchen Tan, Chunwei Lu, Anqi Cui, Han Li, Xi Chen, Kun Xiong, Muzi Wang, Ming Li, Jian Pei, and Jimmy Lin. 2019. Detecting Customer Complaint Escalation with Recurrent Neural Networks and Manually-Engineered Features. In *NAACL-HLT 2019*. 56–63.
[52] Zhitong Yang, Bo Wang, Jinfeng Zhou, Yue Tan, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. 2022. TopKG: Target-oriented Dialog via Global Planning on Knowledge Graph. In *COLING 2022*. 745–755.
[53] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. ReAct: Synergizing Reasoning and Acting in Language Models. *CoRR* abs/2210.03629 (2022).
[54] Weiran Yao, Shelby Heinecke, Juan Carlos Niebles, Zhiwei Liu, Yihao Feng, Le Xue, Rithesh Murthy, Zeyuan Chen, Jianguo Zhang, Devansh Arpit, Ran Xu, Phil Mui, Huan Wang, Caiming Xiong, and Silvio Savarese. 2023. Retroformer: Retrospective Large Language Agents with Policy Gradient Optimization. arXiv:2308.02151 [cs.CL]
[55] Tom Young, Frank Xing, Vlad Pandelea, Jinjie Ni, and Erik Cambria. 2022. Fusing Task-Oriented and Open-Domain Dialogues in Conversational Agents. In *AAAI 2022*. 11622–11629.
[56] Xiao Yu, Maximillian Chen, and Zhou Yu. 2023. Prompt-Based Monte-Carlo Tree Search for Goal-Oriented Dialogue Policy Planning. *CoRR* abs/2305.13660 (2023).
[57] Hamed Zamani, Susan T. Dumais, Nick Craswell, Paul N. Bennett, and Gord Lueck. 2020. Generating Clarifying Questions for Information Retrieval. In *WWW 2020*. 418–428.
[58] Hamed Zamani, Johanne R. Trippas, Jeff Dalton, and Filip Radlinski. 2022. Conversational Information Seeking. *CoRR* abs/2201.08808 (2022).
[59] 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 ACL: EMNLP 2021*. 1092–1101.
[60] Qiang Zhang, Jason Naradowsky, and Yusuke Miyao. 2023. Ask an Expert: Leveraging Language Models to Improve Strategic Reasoning in Goal-Oriented Dialogue Models. In *Findings of ACL: ACL 2023*. 6665–6694.
[61] Shuo Zhang and Krisztian Balog. 2020. Evaluating Conversational Recommender Systems via User Simulation. In *KDD 2020*. 1512–1520.
[62] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. 2018. Towards Conversational Search and Recommendation: System Ask, User Respond. In *CIKM 2018*. 177–186.
[63] Zhonghua Zheng, Lizi Liao, Yang Deng, and Liqiang Nie. 2023. Building Emotional Support Chatbots in the Era of LLMs. *CoRR* abs/2308.11584 (2023).
[64] Yiheng Zhou, Yulia Tsvetkov, Alan W. Black, and Zhou Yu. 2020. Augmenting Non-Collaborative Dialog Systems with Explicit Semantic and Strategic Dialog History. In *ICLR 2020*.