

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

8-2023

A survey on proactive dialogue systems: Problems, methods, and prospects

Yang DENG

Singapore Management University, ydeng@smu.edu.sg

Wenqiang LEI

Wai LAM

Tat-Seng CHUA

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



Part of the [Databases and Information Systems Commons](#), and the [Programming Languages and Compilers Commons](#)

Citation

DENG, Yang; LEI, Wenqiang; LAM, Wai; and CHUA, Tat-Seng. A survey on proactive dialogue systems: Problems, methods, and prospects. (2023). *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, Macao, 2023 August 19-25*. 6583-6591.

Available at: https://ink.library.smu.edu.sg/sis_research/9123

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.

A Survey on Proactive Dialogue Systems: Problems, Methods, and Prospects

Yang Deng¹, Wenqiang Lei^{2,†}, Wai Lam³, Tat-Seng Chua¹

¹National University of Singapore

²Sichuan University

³The Chinese University of Hong Kong

{ydeng,dcscts}@nus.edu.sg, wenqianglei@gmail.com, wlam@se.cuhk.edu.hk

Abstract

Proactive dialogue systems, related to a wide range of real-world conversational applications, equip the conversational agent with the capability of leading the conversation direction towards achieving pre-defined targets or fulfilling certain goals from the system side. It is empowered by advanced techniques to progress to more complicated tasks that require strategic and motivational interactions. In this survey, we provide a comprehensive overview of the prominent problems and advanced designs for conversational agent's proactivity in different types of dialogues. Furthermore, we discuss challenges that meet the real-world application needs but require a greater research focus in the future. We hope that this first survey of proactive dialogue systems can provide the community with a quick access and an overall picture to this practical problem, and stimulate more progresses on conversational AI to the next level.

1 Introduction

Dialogue systems are envisioned to provide social support or functional service to human users via natural language interactions. Conventional dialogue researches mainly focus on the response-ability of the system, such as dialogue context understanding and response generation [Huang *et al.*, 2020]. In terms of the application, typical dialogue systems are designed to passively follow the user-oriented conversation or fulfill the user's request, such as open-domain dialogue systems [Zhang *et al.*, 2018a], task-oriented dialogue systems [Hosseini-Asl *et al.*, 2020], and conversational information-seeking systems [Aliannejadi *et al.*, 2019].

Despite the extensive studies, most dialogue systems typically overlook the design of an essential property in intel-

ligent conversations, *i.e.*, proactivity. Derived from the definition of proactivity in organizational behaviors [Grant and Ashford, 2008] as well as its dictionary definition, conversational agents' proactivity can be defined as the capability to create or control the conversation by taking the initiative and anticipating the impacts on themselves or human users, rather than only passively responding to the users. It will not only largely improve user engagement and service efficiency, but also empower the system to handle more complicated tasks that involve strategic and motivational interactions. Furthermore, the agent's proactivity represents a significant step towards strong AI that has autonomy and human-like consciousness. Even for the powerful ChatGPT, there are still several limitations¹, which attribute to the inability of proactivity, such as passively providing randomly-guessed answers to ambiguous user queries, failing to handle problematic requests that may exhibit harmful or biased conversations.

Several early attempts have been made on enabling the conversational agent to proactively introduce new topics [Li *et al.*, 2016] or useful suggestions [Yan and Zhao, 2018] during the conversation. These pioneering studies have recognized the need for improved problem settings and tangible applications in order to continue enhancing proactive dialogue systems. Recent years have witnessed many advanced designs for conversational agent's proactivity for solving a wide range of challenging dialogue problems. In this survey, we provide a comprehensive review of such efforts that span various task formulations and application scenarios. As shown in Figure 1, we summarize recent studies for three common types of dialogues, namely open-domain dialogues, task-oriented dialogues, and information-seeking dialogues.

Firstly, different from echoing the user's topics, emotions, or views, several problems emerge to enable the system to lead the open-domain dialogues, such as target-guided dialogues [Tang *et al.*, 2019] and prosocial dialogues [Kim *et al.*, 2022]. As the examples illustrated in Figure 2, target-guided dialogues involve the agent leading discussions towards designated target topics (*e.g.*, *Music* to *K-Pop* to *Blackpink*), while prosocial dialogues entrust the agent with constructively guiding conversations according to social norms in response to problematic user utterances (*e.g.*, the cheating intention). Secondly, rather than simply following the user's

* This work was supported by the NEXt Research Centre, the National Natural Science Foundation of China (No. 62272330), the Fundamental Research Funds for the Central Universities of China, the Beijing Academy of Artificial Intelligence (BAAI), and a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14200620).

[†] Corresponding author.

¹ as stated in its official blog <https://openai.com/blog/chatgpt>.

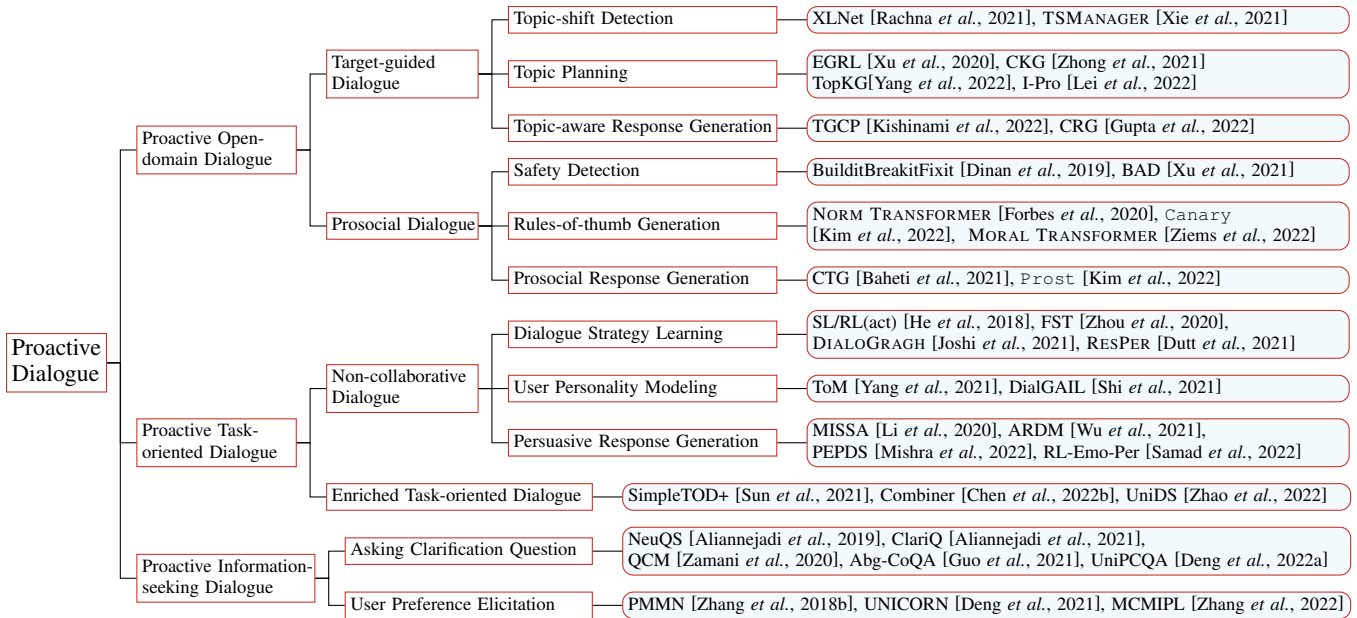


Figure 1: Summary of proactive dialogue systems.

instruction, two distinct types of task-oriented dialogues are characterized by the necessity for the agent’s proactivity: (i) non-collaborative dialogues [Li *et al.*, 2020], where the system and user may have divergent objective or conflicting interests regarding task completion (*e.g.*, the price bargain negotiation), and (ii) enriched task-oriented dialogues [Balaraman and Magnini, 2020], where the agent takes the initiative to provide useful supplementary information not explicitly requested by the user (*e.g.*, additional knowledge or chitchats). Thirdly, we discuss two groups of proactivity designs for enhancing the final performance of conversational information-seeking systems, including asking clarification questions [Aliannejadi *et al.*, 2019] and user preference elicitation [Zhang *et al.*, 2018b]. Accordingly, we introduce the available data resources (summarized in Table 1) and corresponding evaluation protocols for each problem.

In addition, we discuss the main open challenges in developing agent’s proactivity in dialogue systems and several potential research prospects for future studies. (1) Proactivity in Hybrid Dialogues: Hybrid dialogues are the most realistic simulation of interactions between human users and systems, as they incorporate a variety of conversational objectives, instead of focusing a single type of dialogues. Despite the importance of agent’s proactivity in hybrid dialogues, only a few recent studies investigate this critical design. (2) Evaluation Protocols for Proactivity: Compared with general evaluation protocols for dialogue systems, it additionally relies on other disciplines, such as psychology or sociology. Despite this complexity, developing robust and effective evaluation metrics remains critical for advancing techniques in proactive dialogue systems. (3) Ethics of Conversational Agent’s Proactivity: The designs of proactivity in dialogue systems may walk a precarious line between the benefit to human-AI interactions and the potential harm to the human users. Re-

searchers must be mindful of this risk, taking steps to ensure that the systems will not negatively impact the human users.

To our knowledge, this survey is the first to focus on proactive dialogue systems. We systematically summarize recent studies on proactive dialogue systems for three common types of dialogues. We discuss the prominent problems and advanced designs in each type of dialogue as well as present the available data resources and commonly-adopted evaluation protocols. Furthermore, we provide a thorough prospect on several challenging and promising research directions for future studies on proactive dialogue systems.

2 Proactive Open-domain Dialogue Systems

An open-domain dialogue (ODD) system aims to establish long-term connections with users by satisfying the human need for various social supports, such as communication, affection, and belongings [Huang *et al.*, 2020]. In general, the conversational agent is designed to echo the user-oriented topics, emotions, or views. However, such an agent simply acting in a passive role may hinder the conversation progress or introduce undesired issues during the conversation in certain applications. In this section, we present the recent works in terms of two widely-studied problems about proactive ODD systems, *i.e.*, target-guided dialogues [Tang *et al.*, 2019; Wu *et al.*, 2019] and prosocial dialogues [Kim *et al.*, 2022].

2.1 Target-guided Dialogues

Instead of making consistent responses to the user-oriented topics, the dialogue system is required to proactively lead the conversation topics with either a designated target topic [Tang *et al.*, 2019] or a pre-defined topic thread [Wu *et al.*, 2019].

Problem Definition

Given a target t that is only presented to the agent and is unknown to the user, the dialogue begins from an arbitrary

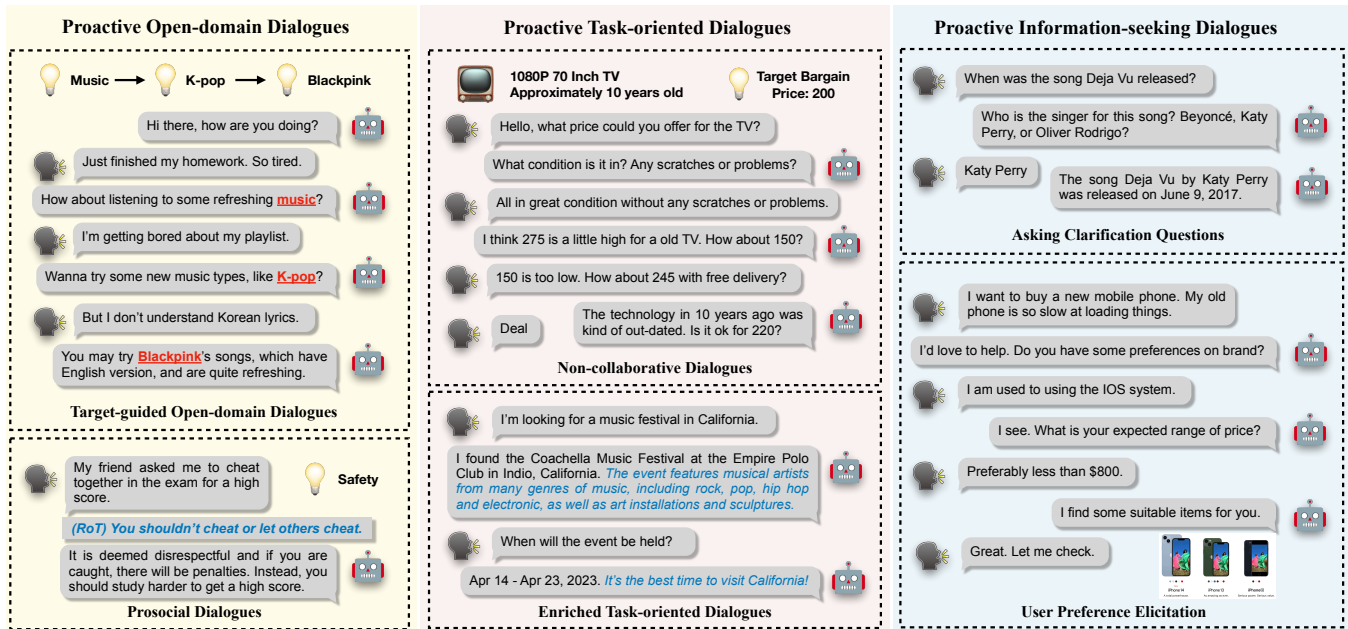


Figure 2: Examples for different problems in proactive dialogue systems.

initial topic, and the system needs to produce multiple turns of responses $\{u_n\}$ to lead the conversation towards the target in the end. The produced responses should satisfy (i) **transition smoothness**, natural and appropriate content under the given dialogue history, and (ii) **target achievement**, driving the conversation to reach the designated target. According to different applications, the target can be a topical keyword [Tang *et al.*, 2019], a knowledge entity [Wu *et al.*, 2019], a conversational goal [Liu *et al.*, 2020], etc. A candidate target set is maintained by the dialogue system.

Methods

There are three main subtasks in target-guided dialogue systems, including topic-shift detection, topic planning, and topic-aware response generation.

- **Topic-shift Detection** aims to promptly discover the topic drift in user utterances. Rachna *et al.* [2021] fine-tune XLNet-base to classify the utterances into major, minor and off topics. Xie *et al.* [2021] construct TIAGE for topic-shift dialogue modeling by augmenting the PersonaChat dataset [Zhang *et al.*, 2018a] with topic-shift annotations, and propose a T5-based topic-shift manager, namely TS-MANAGER, to predict the occurrence of topic shifts.
- **Topic Planning**, which enables the conversation to follow an expected direction, is the core problem in target-guided dialogue systems. Several discourse-level target-guided strategies [Tang *et al.*, 2019; Zhong *et al.*, 2021] constrained on keyword transitions are proposed to proactively drive the conversation topic towards the target. Due to the loose topic-connectivity between keywords, event knowledge graphs are constructed to enhance the coherency in the topic planning [Xu *et al.*, 2020]. However, the knowledge provided in the dialogues is limited for planning a robust

and reasonable topic path towards the target. Therefore, latest studies [Yang *et al.*, 2022] leverage external knowledge graphs for improving the quality of topic transitions with graph reasoning techniques. Instead of corpus-based learning, Lei *et al.* [2022] propose to learn the topic transition from the interactions with users.

- **Topic-aware Response Generation** aims to produce topic-related responses for leading the conversation towards the target. Kishinami *et al.* [2022] propose to generate a complete responding plan that can lead a conversation to the given target. Gupta *et al.* [2022] leverage a bridging path of commonsense knowledge concepts between the current and target topics to generate transition responses.

Datasets and Evaluation Protocols

In this section, we introduce two widely-adopted datasets for the evaluation of target-guided dialogue systems, including Target-Guided Conversation (TGC) [Tang *et al.*, 2019] and DuConv [Wu *et al.*, 2019].

- **TGC** is constructed from Persona-Chat [Zhang *et al.*, 2018a] without the persona information. The target is defined as a keyword in the utterance, which are automatically extracted by a rule-based keyword extractor.
- **DuConv** is constructed by human-human conversations based on two linked entities from the grounded knowledge graph. A grounded knowledge graph is provided for building knowledge-driven proactive dialogue systems.

It is worth noting that Target-Guided Conversation is constructed by labelling targets on existing conversations, while conversely, DuConv is constructed by generating conversations based on designated targets.

Apart from the general evaluation metrics for dialogue systems (BLEU, Dist-N, PPL, etc), we mainly introduce eval-

Dataset	Problem	Language	#Dial.	#Turns	Featured Annotations
TGC [Tang <i>et al.</i> , 2019]	Target-guided Dialogues	English	9,939	11.35	Turn-level Topical Keywords
DuConv [Wu <i>et al.</i> , 2019]	Target-guided Dialogues	Chinese	29,858	9.1	Turn-level Entities & Dialogue-level Goals
MIC [Ziems <i>et al.</i> , 2022]	Prosocial Dialogues	English	38K	2.0	Rules of Thumbs (RoTs) & Revised Responses
ProsocialDialog [Kim <i>et al.</i> , 2022]	Prosocial Dialogues	English	58K	5.7	Safety Labels and Reasons & RoTs
CraigslstBargain [He <i>et al.</i> , 2018]	Non-collaborative Dialogues	English	6,682	9.2	Coarse Dialogue Acts
P4G [Wang <i>et al.</i> , 2019]	Non-collaborative Dialogues	English	1,017	10.43	Dialogue Strategies
ACCENTOR [Sun <i>et al.</i> , 2021]	Enriched Task-oriented Dialogues	English	23.8K	-	Enriched Responses with Chit-chats
KETOD [Chen <i>et al.</i> , 2022b]	Enriched Task-oriented Dialogues	English	5,324	9.78	Turn-level Entities & Enriched Responses with Knowledge
Abg-CoQA [Guo <i>et al.</i> , 2021]	Asking Clarification Questions	English	8,615	5.0	Clarification Need Labels and Questions
PACIFIC [Deng <i>et al.</i> , 2022a]	Asking Clarification Questions	English	2,757	6.89	Clarification Need Labels and Questions

Table 1: Summary of representative data resources that are publicly available for different problems of proactive dialogue systems.

uation protocols that are specific to the concerned problem. The target-guided dialogue systems can be evaluated from two levels, *i.e.*, turn-level and dialogue-level:

- **Turn-level Evaluation:** The performance of target prediction in each turn can be evaluated by (i) $P@K$ and $R@K$, keywords precision and recall at position K in the candidate target set; (ii) Embedding-based correlation scores; and (iii) Proactivity/Smoothness, human evaluation scores that measure how well the system can introduce new topics towards the target while maintaining coherency.
- **Dialogue-level Evaluation:** Due to the high cost and complexity for real user experiments, user simulators are typically adopted for evaluating dialogue-level performance. The common metrics include (i) $SR@t$, success rate of achieving the targets at t -th turn; and (ii) $\#Turns$, the average number of turns used to reach the target.

2.2 Prosocial Dialogues

Most existing dialogue systems fail to handle problematic user utterances by passively agreeing with the unsafe, unethical, or toxic statement, which may cause serious safety concerns for real-world deployment of these systems. Therefore, a trending research topic in proactive dialogue systems is to enable the conversational agent to proactively detect problematic user utterances and constructively and respectfully lead the conversation in a prosocial manner, *i.e.*, following social norms and benefiting others or society [Kim *et al.*, 2022].

Problem Definition

Given a dialogue context, *i.e.*, a sequence of utterances $\{u_1, \dots, u_{t-1}\}$, a prosocial dialogue system aims to first classify the safety label y and then generate a proper response u_t to mitigate the problematic user utterances.

Methods

Current approaches on prosocial dialogue systems can be categorized into three groups, including safety detection, rule-of-thumb generation, and prosocial response generation.

- **Safety Detection.** This line of work focuses on identifying whether the user utterance is problematic for preventing the system generating agreement on problematic statements. Dinan *et al.* [2019] develop a human-in-the-loop training scheme for detecting offensive utterances from other safe utterances in dialogue, which is further improved with adversarial learning [Xu *et al.*, 2021]. Baheti *et al.* [2021] fine-tune offensive language detection classifiers

on a crowd-annotated dataset, TOXICCHAT, which is labeled with offensive language and stance. To avoid classifying specific or sensitive utterances as "unsafe" or "toxic" that may cause social exclusion of minority users, [Kim *et al.*, 2022] introduce a fine-grained safety classification schema: (1) Needs Caution, (2) Needs Intervention, and (3) Casual. Although these approaches can effectively avoid agreeing with problematic user utterances, it is still necessary to continue a user-engaged and socially responsible conversation.

- **Rule-of-Thumb Generation.** RoTs² are generated to interpret why the statement could be seen as acceptable or problematic. Forbes *et al.* [2020] first present SOCIAL-CHEM01, a large-scale corpus for RoT generation, and propose NORM TRANSFORMER to reason about social norms towards the given context. Ziems *et al.* [2022] propose MORAL TRANSFORMER for fine-tuning language models to generate new RoTs that reasonably describe previously unseen dialogue interactions. Kim *et al.* [2022] propose a sequence-to-sequence model `Canary` that generates both safety label and relevant RoTs given a potentially problematic dialogue context. However, the generated RoTs may not be a proper response for user-engaged conversations.
- **Prosocial Response Generation.** Another line of work aims to teach the conversational agent to proactively generate prosocial responses for handling the problematic user utterance. Baheti *et al.* [2021] investigate controllable text generation methods to mitigate the tendency of generating responses that agree with offensive user utterances. Kim *et al.* [2022] propose `Prost` to generate prosocial responses conditioned on the relevant RoTs and the dialogue context.

Datasets and Evaluation Protocols

Many valuable data resources have been constructed for investigating the safety issue in dialogue systems, where most of them focus on preventing the system from generating toxic responses, such as TOXICCHAT [Baheti *et al.*, 2021]. Here we introduce two datasets that provide a more comprehensive evaluation for proactive dialogues, *i.e.*, MORAL INTEGRITY CONVERSATION (MIC) [Ziems *et al.*, 2022] and Prosocial-Dialog [Kim *et al.*, 2022].

- **MIC:** This corpus is constructed by manually annotating prompt-reply pairs (*i.e.*, an open-ended query and an AI-generated response) with Rule-of-Thumbs from SOCIAL-

²**Rule-of-Thumb** (RoT) means a descriptive cultural norm structured as the judgment of an action [Forbes *et al.*, 2020].

CHEM01 [Forbes *et al.*, 2020]. Each RoT serves as a moral judgement that can enhance the original reply.

- **ProsocialDialog:** This corpus is constructed by a human-AI collaboration framework, where AI plays the problematic user role, and crowdworkers play the prosocial agent role, to produce prosocial conversations together. It includes (1) safety labels, (2) RoTs for problematic dialogue contexts, and (3) prosocial responses grounded on RoTs.

According to the three groups of approaches, there are corresponding evaluation metrics: (1) As for safety detection that is intrinsically a classification problem, *Accuracy* and *F1* scores are adopted for evaluation; (2) As for RoT generation and prosocial generation, general text generation metrics (ROUGE, BLEU, PPL) are adopted for evaluation; and (3) Due to the difficulty of measuring prosociality or morality, human evaluation or trained classification models are typically adopted for quantifying different attributes of the generated responses [Ziems *et al.*, 2022; Kim *et al.*, 2022], such as agreement, respect, fairness, etc.

3 Proactive Task-oriented Dialogue Systems

Different from providing social support to users in open-domain dialogues, task-oriented dialogue (TOD) systems target at accomplishing user-requested tasks that are typically domain dependent, such as making reservations or booking tickets. General TOD systems, which typically serve as an obedient assistant to follow the user’s instruction, have achieved promising performance on completing collaborative tasks. Rather than the focus on task completion, there are increasing demands for proactive TOD systems that can either (i) handle non-collaborative tasks where users and systems do not share the same goal, such as negotiation and persuasion dialogues [He *et al.*, 2018; Wang *et al.*, 2019] or (ii) enrich the conversation by providing additional information that is not requested by but useful to the users [Balaraman and Magnini, 2020]. In this section, we introduce the recent advances on these two kinds of proactive TOD systems.

3.1 Non-collaborative Dialogues

Under non-collaborative setting, the system and the user have competing interests or goals towards the task completion but aim to reach an agreement [Zhou *et al.*, 2020]. Typical applications of non-collaborative dialogue systems include negotiating a product price [He *et al.*, 2018], persuading users to make a donation [Wang *et al.*, 2019], deceiving attackers [Li *et al.*, 2020], etc. General TOD systems may fail to meet the system’s goal in these applications due to the only aim of accomplishing the user’s goal. To remedy this issue, the conversational agent needs to successfully develop strategies for conflict resolution and have persuasive powers that can be used to steer the dialogue in a particular direction.

Problem Definition

Given the dialogue history, *i.e.*, a sequence of utterances $\{u_1, \dots, u_{t-1}\}$, along with the previous dialogue strategy sequence $\{s_1, \dots, s_{t-1}\}$ and the dialogue background c , the goal is to generate a response u_t with appropriate dialogue strategy s_t that can lead to a consensus state between the system’s and the user’s goal. Based on different applications, the

dialogue strategy can be coarse dialogue act labels or fine-grained strategy labels, while the dialogue background can be like item descriptions in bargain negotiation or user profile in persuasion dialogues.

Methods

We categorize existing techniques for tackling the non-collaborative dialogue problem into three groups, including dialogue strategy learning, user personality modeling, and persuasive response generation.

- **Dialogue Strategy Learning.** Different from the intent detection in general TOD systems that commonly classify user utterances into pre-defined intents, it further requires the capability of strategic reasoning to handle more complex user actions in non-collaborative dialogues. He *et al.* [2018] aim to control the dialogue strategy to achieve different negotiation goals with the same language generator by decoupling strategy and generation. Zhou *et al.* [2020] employ finite state transducers (FSTs) to leverage effective sequences of strategies in the dialogue context to predict the next strategy. Based on similar motivations, several advanced models have been developed for strategy learning in non-collaborative dialogues, such as DIALOGRAPH [Joshi *et al.*, 2021] with interpretable strategy-graph networks, RESPER [Dutt *et al.*, 2021] with resisting strategy modeling. The strategy learning is still challenging in non-collaborative dialogues, since it involves not only language but also psychological or sociological skills to build rapport and trust between the system and the user.
- **User Personality Modeling.** On the other side, a proactive conversational agent is further required to understand the human decision-making process, where user personality modeling is an important technique. Yang *et al.* [2021] propose to generate strategic dialogue by modeling and inferring personality types of opponents based on the idea of Theory of Mind (ToM) from cognitive science. Shi *et al.* [2021] develop DialGAIL, an RL-based generative algorithm with separate user and system profile builders, to reduce repetition and inconsistency in persuasion dialogues.
- **Persuasive Response Generation.** Since the aim of non-collaborative dialogues is to reach consensus in the end, the responses generated by the system should be persuasive and effective to lead the conversational direction. Following general TOD frameworks, researchers develop modularized [He *et al.*, 2018] and end-to-end [Li *et al.*, 2020; Wu *et al.*, 2021] methods to incorporate persuasive dialogue strategies into response generation. Furthermore, recent studies propose to build empathetic connections between the system and the user for better generating persuasive responses [Mishra *et al.*, 2022; Samad *et al.*, 2022].

Datasets and Evaluation Protocols

Due to the extensive applications of non-collaborative dialogues, many valuable data resources have recently been constructed [Li *et al.*, 2020]. Here we detailedly introduce two widely-adopted benchmarks for evaluating non-collaborative dialogue systems, including CraigslistBargain [He *et al.*, 2018] and PERSUASIONFORGOOD [Wang *et al.*, 2019]. Both of them are collected by human-human conversations.

- **CraigslistBargain**: Two workers are assigned the role of a buyer and a seller and then asked to negotiate the price of an item for sale given a description and photos.
- **PERSUASIONFORGOOD (P4G)** contains persuasion conversations for charity donation, and the corresponding user profiles. It also provides manual annotation of persuasion strategies and dialog acts for each sentence.

As for the evaluation protocols, a key element in non-collaborative dialogues is the strategy learning. Therefore, the accuracy of dialogue strategy prediction is the most commonly-adopted metrics for evaluating the system, such as Accuracy, F1, and ROC AUC scores. For the response generation, some specific aspects are included for human evaluation, such as persuasiveness, task success, etc.

3.2 Enriched Task-oriented Dialogues

The proactivity in general task-oriented dialogue systems commonly refers to the ability of automatically providing additional information that is not requested by but useful to the user [Balaraman and Magnini, 2020], which can improve the quality and effectiveness of conveying functional service in the conversation. The problem formulation of enriched TODs exactly follows that of general TODs, where the difference is that the generated responses in enriched TODs should be not only functionally accurate but also socially engaging. For instance, Sun *et al.* [2021] construct the **ACCENTOR** dataset by adding topical chit-chats into the responses for TODs to make the interactions more engaging and interactive. An end-to-end TOD method, SimpleTOD [Hosseini-Asl *et al.*, 2020], is extended to be SimpleTOD+ for handling enriched TODs, which introduces a new dialogue action, *i.e.*, chit-chat and is further trained on chit-chat generation data. Similarly, [Zhao *et al.*, 2022] develop an end-to-end method, namely UniDS, with a unified dialogue data schema, compatible for both chit-chat and task-oriented dialogues. However, the chit-chat augmentation in **ACCENTOR** mainly contains general greeting responses with limited useful information (*e.g.*, “you’re welcome”). To enrich task-oriented dialogues with knowledgeable chit-chats, Chen *et al.* [2022b] further propose the **KETOD** dataset to enable knowledge-grounded chit-chat regarding relevant entities. A pipeline-based method, namely Combiner, is proposed to reduce the interference between the dialogue state tracking and the generation of knowledge-enriched responses.

4 Proactive Conversational Information Seeking Systems

The goal of conversational information-seeking (CIS) systems is to fulfill the user’s information needs. The typical applications include conversational search, conversational recommendation, and conversational question answering. Conventional CIS systems passively respond to user queries, which may fall short of performing complicated information seeks. Recent years have witnessed several advances on developing proactive CIS systems that can further 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 [Aliannejadi *et al.*, 2021] and conversation question answering [Guo *et al.*, 2021], or elicit the user preference in conversational recommendation [Zhang *et al.*, 2018b].

4.1 Asking Clarification Questions

Asking clarification questions aims to clarify the potential ambiguity in the user query, since the user query is often succinct and brief in real-world conversational search and question answering. The problem is formally formulated by two subtasks [Aliannejadi *et al.*, 2021]: clarification need prediction and clarification question generation. Clarification need prediction is typically viewed as a binary classification problem for predicting whether the user query is ambiguous. If needed, clarification questions can be either selected from a question bank [Aliannejadi *et al.*, 2019] or generated on the fly [Zamani *et al.*, 2020].

Specifically, Aliannejadi *et al.* [2019] propose a question retrieval-selection pipeline, namely NeuQS, to first retrieve top k questions from the question bank and then select the most appropriate question by reranking via BERT-based models. Zamani *et al.* [2020] develop a reinforcement learning based method, namely QCM, to generate clarifying questions by maximizing a clarification utility function. However, these works only focus on the subtask of clarification question generation, while the clarification need prediction is equally important in asking clarification questions. To this end, Aliannejadi *et al.* [2021] and Guo *et al.* [2021] present complete pipeline-based systems for asking clarification questions, which adopt a binary classification model to predict the clarification need label first and then perform clarification question generation. Furthermore, Deng *et al.* [2022a] propose an end-to-end framework, namely UniPCQA, which leverages a unified sequence-to-sequence formulation to tackle three tasks in one model, including clarification need prediction, clarification question generation, and conversational question answering.

Several data resources have recently been constructed for evaluating the system capability of asking clarifying questions for both conversational search, such as Qulac [Aliannejadi *et al.*, 2019], ClariQ [Aliannejadi *et al.*, 2021], and conversational question answering, such as Abg-CoQA [Guo *et al.*, 2021], PACIFIC [Deng *et al.*, 2022a].

4.2 User Preference Elicitation

Instead of simply learning user preference from the dialogue context [Li *et al.*, 2018], Zhang *et al.* [2018b] propose a proactive paradigm, namely “System Ask, User Respond”, to explicitly acquire user preference via asking questions in conversational recommendation. The problem is formulated as predicting the item attribute for eliciting user preferences at the next turn, *e.g.*, “Which brand of laptop do you prefer?”.

A personalized multi-memory network (PMMN) [Zhang *et al.*, 2018b] is first designed to incorporate user embeddings into next question prediction at turn-level. Due to the complexity of user preferences, multiple turns of question asking are required. Therefore, recent works tackle the user preference elicitation at dialogue-level, *i.e.*, “what questions to

ask”, as a multi-step decision making process by reinforcement learning (RL) [Deng *et al.*, 2021; Zhang *et al.*, 2022]. Deng *et al.* [2021] propose a graph-based RL framework for policy learning, namely UNICORN, which models real-time user preference during the conversation with a dynamic weighted graph structure. Motivated by the complex user interests in CRS, Zhang *et al.* [2022] propose the MCMIPL framework to efficiently obtain user preferences by asking multi-choice questions.

Since the goal of CRS is to make successful recommendations, its evaluation is basically based on the final recommendation results. Similar to target-guided dialogues in Section 2.1, the evaluation protocols of question-driven CRS can also be divided into turn-level and dialogue-level:

- **Turn-level:** The hit ratio ($HR@k, t$) is adopted to evaluate the next question prediction regarding the top- k predicted attributes at conversation turn t . General recommendation metrics, such as $MRR@k, t$, $MAP@k, t$, $NDCG@k, t$, are adopted to evaluate the item recommendation regarding the top- k ranked items based on the elicited user preference.
- **Dialogue-level:** The success rate at the turn t ($SR@t$) is adopted to measure the cumulative ratio of successful recommendation by the turn t . AT is adopted to evaluate the average number of turns for all sessions.

Despite the advances on context-driven CRS [Li *et al.*, 2018] and question-driven CRS [Zhang *et al.*, 2018b] separately, it would be beneficial to develop CRS that can effectively combine the advantages of these two kinds of CRS, *i.e.*, natural language understanding and generation as well as proactive user preference elicitation. Meanwhile, current studies on user preference elicitation are basically evaluated on synthetic conversation data from product reviews [Zhang *et al.*, 2018b] or purchase logs [Deng *et al.*, 2021; Zhang *et al.*, 2022]. Therefore, well-constructed benchmarks with human-human conversations are still in great demand for facilitating more robust and reliable evaluations.

5 Challenges and Prospects

In this section, we discuss some challenges that meet real-world application needs but require more research efforts.

- **Proactivity in Hybrid Dialogues.** All the aforementioned conversational systems assume that users always have a clear conversational goal and the system also solely targets at reaching a certain goal, such as chit-chat, question answering, recommendation, etc. The system with a higher level of agent’s proactivity should also be capable of handling conversations with multiple and various goals. Recently, many efforts have been made on constructing valuable data resources for hybrid dialogue systems with multiple conversational goals, such as DuRecDial [Liu *et al.*, 2020], FusedChat [Young *et al.*, 2022], SalesBot [Chiu *et al.*, 2022], and OB-MultiWOZ [Li *et al.*, 2022b]. Early studies simply tackle this problem similar to topic-guided response generation with pre-defined goals [Zhang *et al.*, 2021]. While some latest works [Liu *et al.*, 2022; Deng *et al.*, 2022b] argue the necessity of proactively discovering users’ interests and naturally leading user-engaged dialogues with changing conversational goals. In practice,

hybrid dialogue systems are the closest form of real-world applications. More efforts should be made to ensure natural and smooth transitions among different types of dialogues as well as improve the overall dialogue quality without performance loss of certain types of dialogues.

- **Evaluation Protocols for Proactivity.** The development of robust evaluation protocols has already been a long-standing problem for different kinds of conversational systems. Most studies on proactive dialogue systems adopt human evaluation for manually assessing the dialogue quality. To mitigate the high cost of interacting with real users, building user simulators is a relatively efficient and labourless techniques for evaluating the proactive interactions in dialogue systems [Tang *et al.*, 2019; Sekulic *et al.*, 2022]. Besides, goal completion and user satisfaction are identified as two essential metrics for evaluating proactive conversational systems [Liu *et al.*, 2020; Yang *et al.*, 2022; Lei *et al.*, 2022; Gupta *et al.*, 2022]. In fact, the evaluation for conversational agent’s proactivity is a more challenging problem, since it involves the evaluation not only from the perspective of natural language, but also from the other disciplines, such as human-computer interaction, sociology, psychology, etc. The rapid development of proactive dialogue systems urges more effective and robust multidisciplinary evaluation protocols.
- **Ethics of Conversational Agent’s Proactivity.** Designing agent’s proactivity in dialogue systems is inevitably a double-edged sword that can be used for good or evil. For responsible AI researches, it is necessary to guarantee several important aspects of ethical issues in proactive dialogue systems: (i) **Factuality:** Factual incorrectness and hallucination of knowledge are common in dialogue systems, while the agent’s proactivity will introduce more system-initiative information. Therefore, besides the generation process, it is also crucial to guarantee the factuality of the external knowledge [Chen *et al.*, 2022a]. (ii) **Morality:** Besides general dialogue morality problems, such as toxic language and social bias [Kim *et al.*, 2022; Ziems *et al.*, 2022], proactive dialogue systems also need to pay attention to the aggressiveness issue during the non-collaborative conversations [Mishra *et al.*, 2022; Samad *et al.*, 2022]. (iii) **Privacy:** The privacy issue is overlooked in current studies on dialogue systems [Li *et al.*, 2022a], but the agent’s proactivity raises more concerns on misusing personal information obtained from the users.

6 Conclusion

This paper provides a comprehensive review on the problems and designs of conversational agent’s proactivity in different types of dialogue systems. Accordingly, we summarize the representative data resources that are publicly available for investigating the corresponding problems in proactive dialogue systems. Furthermore, we shed light on several open challenges in this field, ultimately aiming to advance research and innovation for developing proactive dialogue systems. As a significant step towards the next-generation of conversational AI, it necessitates increased research and further exploration in the development of proactive dialogue systems.

References

- [Aliannejadi *et al.*, 2019] Mohammad Aliannejadi, Hamed Zamani, Fabio Crestani, and W. Bruce Croft. Asking clarifying questions in open-domain information-seeking conversations. In *SIGIR*, 2019.
- [Aliannejadi *et al.*, 2021] Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail S. Burtsev. Building and evaluating open-domain dialogue corpora with clarifying questions. In *EMNLP*, 2021.
- [Baheti *et al.*, 2021] Ashutosh Baheti, Maarten Sap, Alan Ritter, and Mark O. Riedl. Just say no: Analyzing the stance of neural dialogue generation in offensive contexts. In *EMNLP*, 2021.
- [Balaraman and Magnini, 2020] Vevake Balaraman and Bernardo Magnini. Proactive systems and influenceable users: Simulating proactivity in task-oriented dialogues. In *SEMDIAL*, 2020.
- [Chen *et al.*, 2022a] Maximillian Chen, Weiyang Shi, Feifan Yan, Ryan Hou, Jingwen Zhang, Saurav Sahay, and Zhou Yu. Seamlessly integrating factual information and social content with persuasive dialogue. In *AAACL/IJCNLP*, 2022.
- [Chen *et al.*, 2022b] Zhiyu Chen, Bing Liu, Seungwhan Moon, Chinnadhurai Sankar, Paul A. Crook, and William Yang Wang. KETOD: knowledge-enriched task-oriented dialogue. In *Findings of ACL: NAACL*, 2022.
- [Chiu *et al.*, 2022] Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. Salesbot: Transitioning from chat to task-oriented dialogues. In *ACL*, 2022.
- [Deng *et al.*, 2021] Yang Deng, Yaliang Li, Fei Sun, Bolin Ding, and Wai Lam. Unified conversational recommendation policy learning via graph-based reinforcement learning. In *SIGIR*, 2021.
- [Deng *et al.*, 2022a] Yang Deng, Wenqiang Lei, Wenxuan Zhang, Wai Lam, and Tat-Seng Chua. PACIFIC: towards proactive conversational question answering over tabular and textual data in finance. In *EMNLP*, 2022.
- [Deng *et al.*, 2022b] Yang Deng, Wenxuan Zhang, Weiwen Xu, Wenqiang Lei, Tat-Seng Chua, and Wai Lam. A unified multi-task learning framework for multi-goal conversational recommender systems. *TOIS*, 2022.
- [Dinan *et al.*, 2019] Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. Build it break it fix it for dialogue safety: Robustness from adversarial human attack. In *EMNLP-IJCNLP*, 2019.
- [Dutt *et al.*, 2021] Ritam Dutt, Sayan Sinha, Rishabh Joshi, Surya Shekhar Chakraborty, Meredith Riggs, Xinru Yan, Haogang Bao, and Carolyn P. Rosé. Resper: Computationally modelling resisting strategies in persuasive conversations. In *EACL*, 2021.
- [Forbes *et al.*, 2020] Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. Social chemistry 101: Learning to reason about social and moral norms. In *EMNLP*, 2020.
- [Grant and Ashford, 2008] Adam M Grant and Susan J Ashford. The dynamics of proactivity at work. *Research in organizational behavior*, 2008.
- [Guo *et al.*, 2021] Meiqi Guo, Mingda Zhang, Siva Reddy, and Malihe Alikhani. Abg-coqa: Clarifying ambiguity in conversational question answering. In *AKBC*, 2021.
- [Gupta *et al.*, 2022] Prakhar Gupta, Harsh Jhamtani, and Jeffrey P. Bigham. Target-guided dialogue response generation using commonsense and data augmentation. In *Findings of ACL: NAACL*, 2022.
- [He *et al.*, 2018] He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. Decoupling strategy and generation in negotiation dialogues. In *EMNLP*, 2018.
- [Hosseini-Asl *et al.*, 2020] Ehsan Hosseini-Asl, Bryan McCann, Chien-Sheng Wu, Semih Yavuz, and Richard Socher. A simple language model for task-oriented dialogue. In *NeurIPS*, 2020.
- [Huang *et al.*, 2020] Minlie Huang, Xiaoyan Zhu, and Jianfeng Gao. Challenges in building intelligent open-domain dialog systems. *ACM Trans. Inf. Syst.*, 2020.
- [Joshi *et al.*, 2021] Rishabh Joshi, Vidhisha Balachandran, Shikhar Vashishth, Alan W. Black, and Yulia Tsvetkov. Dialograph: Incorporating interpretable strategy-graph networks into negotiation dialogues. In *ICLR*, 2021.
- [Kim *et al.*, 2022] Hyunwoo Kim, Youngjae Yu, Liwei Jiang, Ximing Lu, Daniel Khashabi, Gunhee Kim, Yejin Choi, and Maarten Sap. Prosocialdialog: A prosocial backbone for conversational agents. In *EMNLP*, 2022.
- [Kishinami *et al.*, 2022] Yosuke Kishinami, Reina Akama, Shiki Sato, Ryoko Tokuhisa, Jun Suzuki, and Kentaro Inui. Target-guided open-domain conversation planning. In *COLING*, 2022.
- [Lei *et al.*, 2022] Wenqiang Lei, Yao Zhang, Feifan Song, Hongru Liang, Jiabin Mao, Jiancheng Lv, Zhenglu Yang, and Tat-Seng Chua. Interacting with non-cooperative user: A new paradigm for proactive dialogue policy. In *SIGIR*, 2022.
- [Li *et al.*, 2016] Xiang Li, Lili Mou, Rui Yan, and Ming Zhang. Stalematebreaker: A proactive content-introducing approach to automatic human-computer conversation. In *IJCAI*, 2016.
- [Li *et al.*, 2018] Raymond Li, Samira Ebrahimi Kahou, Hannes Schulz, Vincent Michalski, Laurent Charlin, and Chris Pal. Towards deep conversational recommendations. In *NeurIPS*, 2018.
- [Li *et al.*, 2020] Yu Li, Kun Qian, Weiyang Shi, and Zhou Yu. End-to-end trainable non-collaborative dialog system. In *AAAI*, 2020.
- [Li *et al.*, 2022a] Haoran Li, Yangqiu Song, and Lixin Fan. You don't know my favorite color: Preventing dialogue representations from revealing speakers' private personas. In *NAACL*, 2022.
- [Li *et al.*, 2022b] Miaoran Li, Baolin Peng, Jianfeng Gao, and Zhu Zhang. OPERA: harmonizing task-oriented dialogues and information seeking experience. *CoRR*, 2022.

- [Liu *et al.*, 2020] Zeming Liu, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. Towards conversational recommendation over multi-type dialogs. In *ACL*, 2020.
- [Liu *et al.*, 2022] Zeming Liu, Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, and Hua Wu. Where to go for the holidays: Towards mixed-type dialogs for clarification of user goals. In *ACL*, 2022.
- [Mishra *et al.*, 2022] Kshitij Mishra, Azlaan Mustafa Samad, Palak Totala, and Asif Ekbal. PEPDS: A polite and empathetic persuasive dialogue system for charity donation. In *COLING*, 2022.
- [Rachna *et al.*, 2021] Konigari Rachna, Saurabh Ramola, Vijay Vardhan Alluri, and Manish Shrivastava. Topic shift detection for mixed initiative response. In *SIGdial*, 2021.
- [Samad *et al.*, 2022] Azlaan Mustafa Samad, Kshitij Mishra, Mauajama Firdaus, and Asif Ekbal. Empathetic persuasion: Reinforcing empathy and persuasiveness in dialogue systems. In *Findings of ACL: NAACL*, 2022.
- [Sekulic *et al.*, 2022] Ivan Sekulic, Mohammad Aliannejadi, and Fabio Crestani. Evaluating mixed-initiative conversational search systems via user simulation. In *WSDM*, 2022.
- [Shi *et al.*, 2021] Weiyang Shi, Yu Li, Saurav Sahay, and Zhou Yu. Refine and imitate: Reducing repetition and inconsistency in persuasion dialogues via reinforcement learning and human demonstration. In *Findings of ACL: EMNLP*, 2021.
- [Sun *et al.*, 2021] Kai Sun, Seungwhan Moon, Paul A. Crook, Stephen Roller, Becka Silvert, Bing Liu, Zhiguang Wang, Honglei Liu, Eunjoon Cho, and Claire Cardie. Adding chit-chat to enhance task-oriented dialogues. In *NAACL-HLT*, 2021.
- [Tang *et al.*, 2019] Jianheng Tang, Tiancheng Zhao, Chenyan Xiong, Xiaodan Liang, Eric P. Xing, and Zhiting Hu. Target-guided open-domain conversation. In *ACL*, 2019.
- [Wang *et al.*, 2019] Xuewei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. Persuasion for good: Towards a personalized persuasive dialogue system for social good. In *ACL*, 2019.
- [Wu *et al.*, 2019] Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. Proactive human-machine conversation with explicit conversation goal. In *ACL*, 2019.
- [Wu *et al.*, 2021] Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu. Alternating recurrent dialog model with large-scale pre-trained language models. In *EACL*, 2021.
- [Xie *et al.*, 2021] Huiyuan Xie, Zhenghao Liu, Chenyan Xiong, Zhiyuan Liu, and Ann A. Copestake. TIAGE: A benchmark for topic-shift aware dialog modeling. In *Findings of ACL: EMNLP*, 2021.
- [Xu *et al.*, 2020] Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu, and Wanxiang Che. Enhancing dialog coherence with event graph grounded content planning. In *IJCAI*, 2020.
- [Xu *et al.*, 2021] Jing Xu, Da Ju, Margaret Li, Y. Boureau, Jason Weston, and Emily Dinan. Bot-adversarial dialogue for safe conversational agents. In *NAACL-HLT*, 2021.
- [Yan and Zhao, 2018] Rui Yan and Dongyan Zhao. Smarter response with proactive suggestion: A new generative neural conversation paradigm. In *IJCAI*, 2018.
- [Yang *et al.*, 2021] Runzhe Yang, Jingxiao Chen, and Karthik Narasimhan. Improving dialog systems for negotiation with personality modeling. In *ACL/IJCNLP*, 2021.
- [Yang *et al.*, 2022] Zhitong Yang, Bo Wang, Jinfeng Zhou, Yue Tan, Dongming Zhao, Kun Huang, Ruifang He, and Yuexian Hou. Topkg: Target-oriented dialog via global planning on knowledge graph. In *COLING*, 2022.
- [Young *et al.*, 2022] Tom Young, Frank Xing, Vlad Pandlea, Jinjie Ni, and Erik Cambria. Fusing task-oriented and open-domain dialogues in conversational agents. In *AAAI*, 2022.
- [Zamani *et al.*, 2020] Hamed Zamani, Susan T. Dumais, Nick Craswell, Paul N. Bennett, and Gord Lueck. Generating clarifying questions for information retrieval. In *WWW*, 2020.
- [Zhang *et al.*, 2018a] Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing dialogue agents: I have a dog, do you have pets too? In *ACL*, 2018.
- [Zhang *et al.*, 2018b] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W. Bruce Croft. Towards conversational search and recommendation: System ask, user respond. In *CIKM*, 2018.
- [Zhang *et al.*, 2021] Jun Zhang, Yan Yang, Chencai Chen, Liang He, and Zhou Yu. KERS: A knowledge-enhanced framework for recommendation dialog systems with multiple subgoals. In *Findings of ACL: EMNLP*, 2021.
- [Zhang *et al.*, 2022] Yiming Zhang, Lingfei Wu, Qi Shen, Yitong Pang, Zhihua Wei, Fangli Xu, Bo Long, and Jian Pei. Multiple choice questions based multi-interest policy learning for conversational recommendation. In *WWW*, 2022.
- [Zhao *et al.*, 2022] Xinyan Zhao, Bin He, Yasheng Wang, Yitong Li, Fei Mi, Yajiao Liu, Xin Jiang, Qun Liu, and Huanhuan Chen. Unids: A unified dialogue system for chit-chat and task-oriented dialogues. In *DialDoc@ACL*, 2022.
- [Zhong *et al.*, 2021] Peixiang Zhong, Yong Liu, Hao Wang, and Chunyan Miao. Keyword-guided neural conversational model. In *AAAI*, 2021.
- [Zhou *et al.*, 2020] Yiheng Zhou, Yulia Tsvetkov, Alan W. Black, and Zhou Yu. Augmenting non-collaborative dialog systems with explicit semantic and strategic dialog history. In *ICLR*, 2020.
- [Ziems *et al.*, 2022] Caleb Ziems, Jane A. Yu, Yi-Chia Wang, Alon Y. Halevy, and Diyi Yang. The moral integrity corpus: A benchmark for ethical dialogue systems. In *ACL*, 2022.